DIVERSE DISCIPLINES, ONE COMMUNITY

Biomedical Computation REVIEW

Published by the Mobilize Center, an NIH Big Data to Knowledge Center of Excellence

Wearing Your Health *on Your Sleeve:*

How big data from mobile apps and sensors may revolutionize healthcare

PLUS:
Integrating the Fragmented Mind

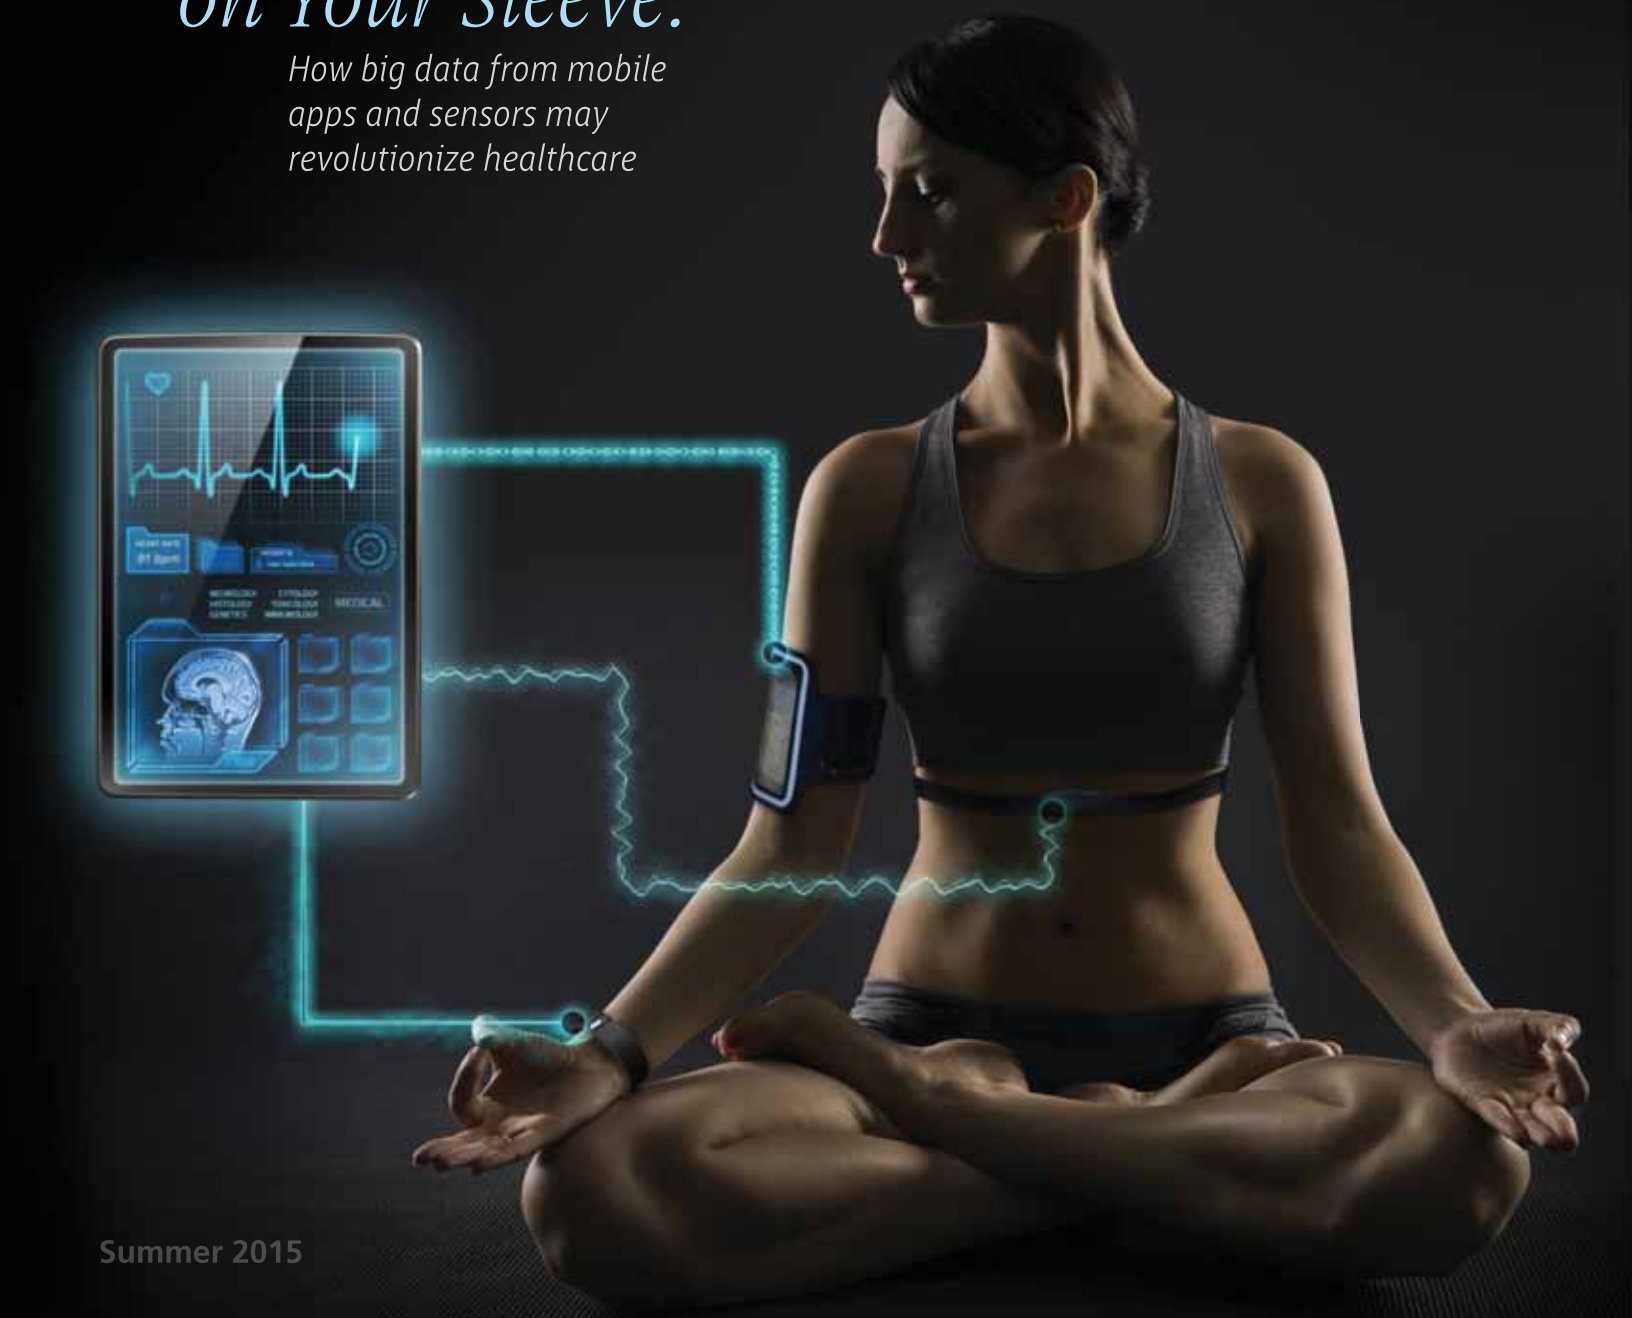

Summer 2015

contents

ContentsSummer 2015

DEPARTMENTS

Cover and Page 14 Art: Created by Rachel Jones of Wink Design Studio using © Adamgregor | Dreamstime.com (heart rate monitor), © Štěpán Kápl | Dreamstime.com (modern healthcare device), © Vera Volkova | Dreamstime.com (armband), © Inara Prusakova | Dreamstime.com (yoga woman sitting in lotus pose), and Fitbit image.

Page 22 Art: Created by Rachel Jones of Wink Design Studio using © Antares614 | Dreamstime.com (asian fan image), © Liouthe | Dreamstime.com (spear), © Blackslide | Dreamstime.com (rope), © Yakobchuk | Dreamstime.com (brain), and © Carolyne Pehora | Dreamstime.com (elephant). Other photo and texture elements by Wink Design Studio.

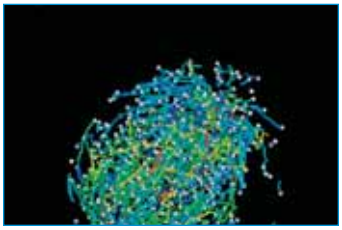

Summer 2015
Volume 11, Issue 2
ISSN 1557-3192

For general inquiries, subscriptions, or letters to the editor,
visit our website at www.bcr.org

Office
Biomedical Computation Review
Stanford University
318 Campus Drive
Clark Center Room S271
Stanford, CA 94305-5444

Biomedical Computation Review
is published by:

mobilize
The Mobilize Center, an NIH
Big Data to Knowledge (BD2K)
Center of Excellence
mobilize.stanford.edu

Publication is supported by NIH Big Data to Knowledge (BD2K) Research Grant U54EB020405. Information on the BD2K program can be found at http://datascience.nih.gov/bd2k. The NIH program and science officers for the Mobilize Center are:

Grace Peng, National Institute of Biomedical Imaging and Bioengineering,

Theresa Cruz, National Institute of Child Health and Human Development, and

Daofen Chen, National Institute of Neurological Disorders and Stroke

GuestEditorial

BY JOSEPH C. KVEDAR, MD, VICE PRESIDENT, CONNECTED HEALTH, PARTNERS HEALTHCARE

Wearable Technology: Making Health Addictive

The hype around wearables is deafening. I say this from the perspective of someone who saw their application in chronic illness management more than 15 years ago. Of course, at that time, it was less about wearables and more about sensors in the home, but the concept was the same.

Partners HealthCare has been committed to technology-enabled care for nearly two decades. Over the years, we've seen growing signs that wearables were going to be all the rage. In 2005, we coined the term 'Connected Health' and the slogan, "Bring health care into the day-to-day lives of our patients," shortly thereafter. About 18 months ago, Partners Connected Health launched Wellocracy, an online community designed to educate consumers about the power of self-tracking as a tool for health improvement. All of this attention to wearables warms my heart.

Sarah C.P. Williams has written an informative feature story in this issue that clearly defines what is required to make these technologies applicable in healthcare and what to do with all of the patient data being collected. Based on our clinical research at Partners, including connected health programs in chronic disease management, adherence and wellness, we have verified many of the principles mentioned in this article.

To add further perspective, there are a few additional concepts that I believe we must get right in order to harness mobile technology. Research has shown that people check their mobile phones upwards of 150 times per day. So how do we leverage mobile health to make health and wellness as addictive?

Making health addictive is really about harnessing the power of our fascination with mobile devices, particularly smartphones. Could we induce permanent behavior change if we put a personalized, relevant, motivational and unobtrusive message in front of you some of those many times you check your mobile device?

Today, most healthcare app development is still confusing education with inspiration. We are learning a great deal about how to empower patients to self-manage their health, and what to do with all of this patient-generated data. The one critical element we must get right is how to 'sell' health to consumers and keep them coming back for more. Again, it's got to be personal, motivational and ubiquitous.

Research has shown that people check their mobile phones upwards of 150 times per day. So how do we leverage mobile health to make health and wellness as addictive?

Making health addictive is really about harnessing the power of our fascination with mobile devices, particularly smartphones.

Looking ahead, we must ask the question, is the future of patient-generated data migrating to the mobile phone or will it move into the realm of micro-sized wearable seeds, ingestibles, injectables, bandaids and the like? There are also many more health sensing applications than just pure activity tracking, such as continuous heart rate or blood pressure monitoring.

The power of sensor-generated data in personal health and chronic illness management is simply too powerful to ignore. □

BY KATHARINE MILLER

Giving Away the NLP Store

In order for computers to extract knowledge from the vastness of all biomedical literature, the machines must first determine the structure of the natural language text—what are the nouns and verbs and what is their role in each sentence. But transforming such a large body of literature into a series of dependency trees not unlike the sentence diagrams of old-time grammar classes requires millions of CPU hours. Fortunately, researchers associated with the Mobilize Center have completed this preprocessing task for several large biomedical literature datasets (including PubMed Central's Open Access Subset, PLoS and BioMed Central)—and are offering the marked-up resources for use by others.

"We're really excited that we've been able to preprocess these datasets and make them available," says **Chris Ré, PhD**, assistant professor in the computer science department at Stanford University. "Now other researchers who want to prospect the scientific literature and perform natural language processing can get all that detailed markup for free. We hope they find something interesting with it."

Most labs can't afford to preprocess such a huge amount of literature, Ré says. But in collaboration with the Center for High Throughput Computing at the University of Wisconsin, Ré's group gained access to many hours of computational time using the Open Science Grid. This allowed them to mark up large volumes of creative commons literature for their own work—and now offer it to others.

Before the end of 2015, Ré plans to use DeepDive, a free, open-source, probabilistic inference engine developed by his lab, to go a step further with his natural language processing (NLP) analysis of the biomedical literature and release those results for free as well. If one thinks of NLP preprocessing as identifying the nouns, verbs, and objects (x inhibits y, for example) then DeepDive might be determining the nature of the entities—whether x and y are genes or proteins, for example—as well as the genes' or proteins' relationship to some specific disease term described in the same paper. "These inference or entity resolution problems are very challenging computationally as well as challenging to get high quality on," Ré says. But DeepDive has proven adept at the task, sometimes outperforming expert annotators.

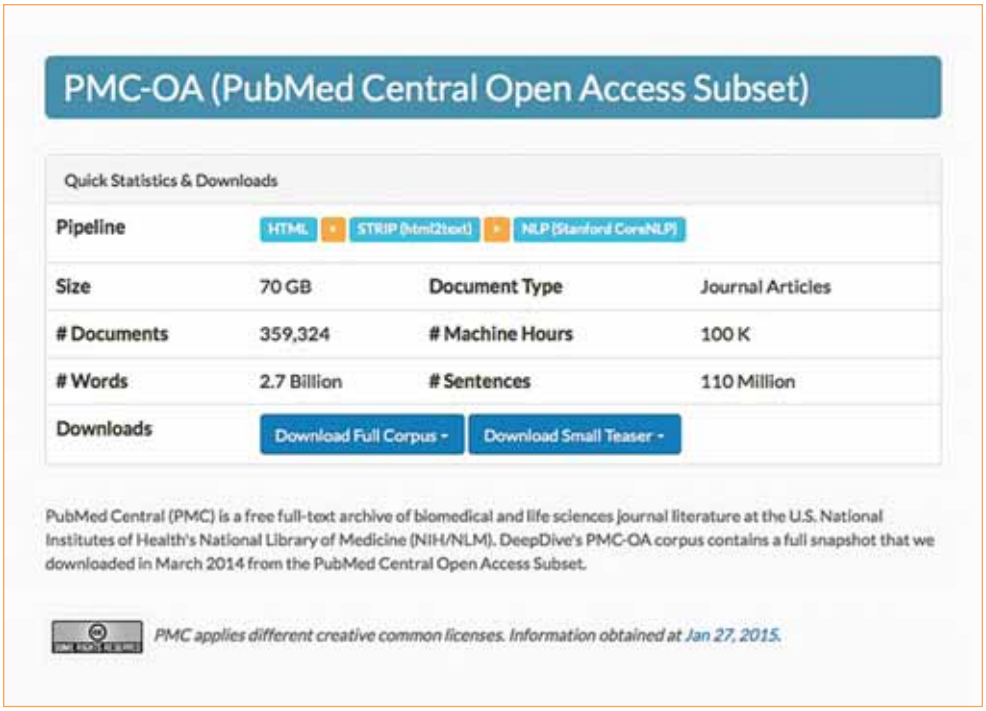

Screenshot from the DeepDive Open Datasets web page offering preprocessed NLP analysis of the open access portion of PubMed Central.

How it works: Domain experts specify the kinds of relationships or features they are interested in. They might provide examples from ontologies or a sample of manually curated data, or they might just explain to DeepDive researchers how to reason with the data. "We take that specification and translate it into a large probabilistic inference problem," Ré says. "We solve that and produce data for the researchers." It's an iterative process—the researchers look at what comes out and give feedback that is used to train the system over time, so it can extract the entities or relationships more robustly.

But the release of the NLP datasets should be useful to people even if they don't use DeepDive, Re says. Right now, people are still just downloading the datasets and poking around. "The excitement is: we're giving away data," Re says. "If it's useful to anyone, send us a note." □

DETAILS

The DeepDive Open Datasets are available at http://deepdive.stanford.edu/doc/opendata/

Big DataHighlight

BY KATHARINE MILLER

111 Reference Human Epigenomes

Each cell in a person's body contains the same three billion letter DNA encyclopedia. But the many different cell types throughout the body (brain, bone, heart, skin, etc.) represent different readers of that encyclopedia who have each highlighted their favorite parts, dog-eared certain pages, annotated interesting paragraphs, and crossed out things they find dull or uninteresting. These markings constitute the epigenome. "The epigenome tells us what are the important parts to read," says **Manolis Kellis, PhD**, professor of computer science at the Massachusetts Institute of Technology (MIT).

And now the marked up encyclopedias are available to all. A February 2015 *Nature* paper by the Roadmap Epigenomics Consortium, of which Kellis is a part, reports that they have mapped the epigenomes for 111 different human cell types. In addition, the researchers compared the different cell types' epigenomic signatures and even uncovered some possible ways the epigenome may play a role in disease.

Epigenetic changes in different kinds of cells include such things as DNA methylation and modifications to the histones around which DNA is wrapped (like beads on a string) in the cell nucleus. These kinds of changes affect which genes in the cell are expressed at any particular time.

To map all such changes for 111 different human cell types—an enormous task in itself—the Consortium used a variety of assays to pull out the modified parts of the genome and then sequence the attached DNA to determine its location in the genome. This produced an epigenomic map showing each kind of modification for each type of cell.

The next step, Kellis says, was to figure out how the marks have meaning. For that, he and his colleagues turned to hidden Markov models capable of identifying underlying patterns in the epigenome maps. "Some patterns happen at the starts of genes; others in the gene; others in places that are repressed," Kellis says. These patterns revealed regulatory modules of coordinated activity as well as their likely activators and repressors.

Once that was understood, Kellis says, "Then we could start studying the differences." And some interesting ones were noted. For example, many cells derived from embryonic stem cells clustered closely with other embryonic stem cells rather than their corresponding tissues, suggesting their stem cell nature was still predominant. For one particular enhancer signal, fetal brain and germinal matrix cells cluster with neural stem cells rather than adult brain cells.

Going a step further, the researchers looked at how the epigenomic data sets squared with disease-associated variants identified in various genome-wide association studies (GWAS). In many cases, disease variants were enriched in epigenomic modifications for trait-specific tissues. For example, the team looked across the genome at all the genetic variants associated with blood pressure and found that they tend to be active in cells in the left ventricle of the heart. "These are genetic differences that affect the circuitry that turns the genes on and off," Kellis says. Thus, the way cells are controlled in the left ventricle has something to do with how much pressure builds up.

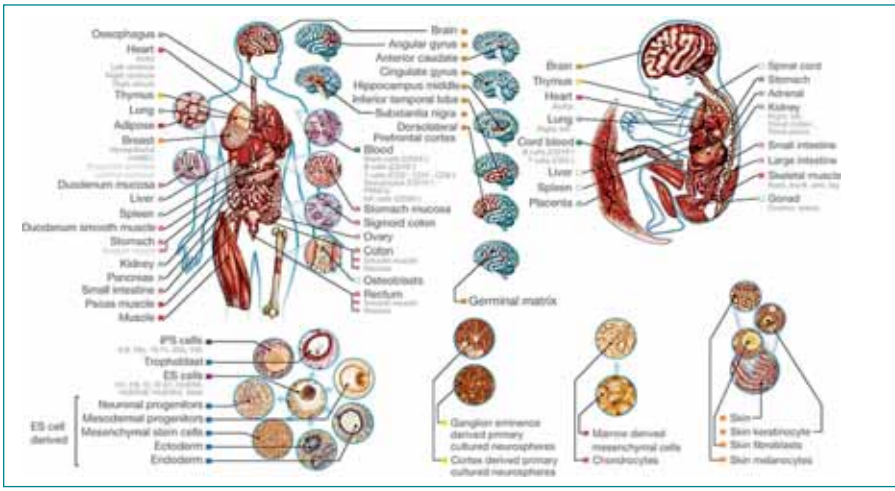

Researchers have sequenced the epigenomes of 111 different human cell types. Reprinted by permission from Macmillan Publishers Ltd: Roadmaps Epigenomic Consortium et al., Integrative analysis of 111 reference human epigenomes,* Nature *518:318 (2015).

A separate *Nature* paper by the Consortium also reported that the epigenomic signature for a cancerous tumor could help identify the originating site of the cancer—a piece of information that is sometimes unknown and confounds appropriate treatment.

This work is important, Kellis says, because it provides access not just to the protein coding part of the genome, but also to the hidden control circuitry. "We can read the circuitry of the genome and understand which genes are active and use that to understand where the genetic predispositions to disease lie within the genome," he says. □

This work is important, Kellis says, because it provides access not just to the protein coding part of the genome, but also to the hidden control circuitry.

SERIOUS GAMES AND BIOMEDICAL RESEARCH: What Do Game Developers Bring to the Table?

By Alexander Gelfand

Some biomedical researchers are serious about games.

"Just because something is a game does not mean it is childish," says **Ingmar Riedel-Kruse, PhD**, assistant professor of bioengineering at Stanford University.

Indeed, the so-called serious games movement has established that games in general—and electronic games in particular—can serve as tools for accomplishing meaningful goals, like helping people to improve their eating habits or understand the federal budget.

For his part, Riedel-Kruse has developed a number of "biotic games" that employ real biological materials. Right now, these are primarily educational in nature, allowing players to influence, observe, and understand the behavior of simple living organisms. In PAC-mecium, for example, players use something resembling a conventional videogame controller to herd a flock of actual paramecia, single-celled organisms that live in ponds and respond readily to mild electrical stimulation. The paramecia are represented on a monitor by a digital image of a fish; and while players can also watch real-time video of the tiny creatures swimming about and changing direction in response to electrical fields, the fun lies in using their digital fish avatar to gobble up virtual yeast pellets while trying to avoid the bite of a predatory zebrafish.

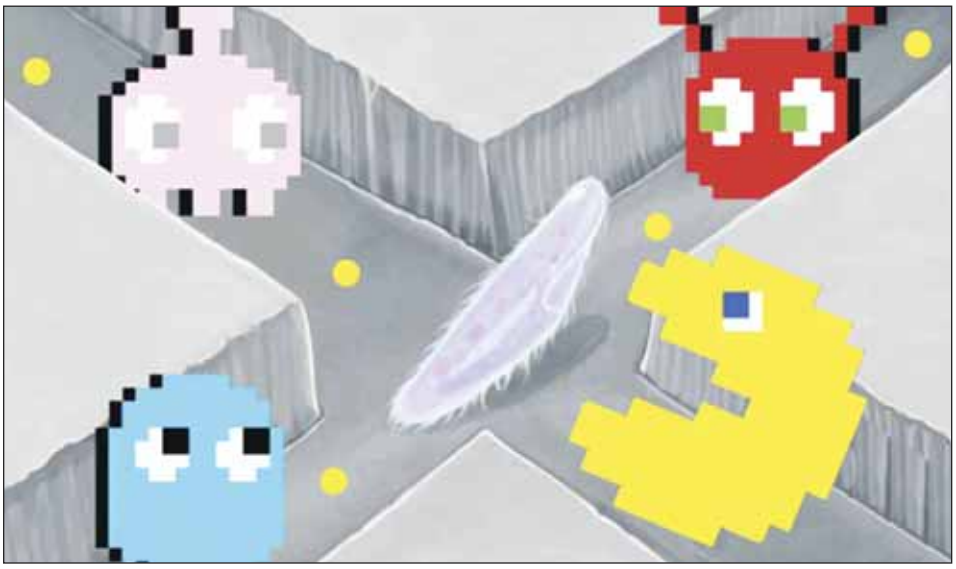

In the biotic game PAC-Mecium, players redirect the movements of actual live paramecia by changing the surrounding electric potential. Reprinted from H. Riedel-Kruse, A. M. Chung, B. Dura, A. L. Hamilton and B. C. Lee, Design, engineering and utility of biotic games, Lab Chip, *11:14-22 (2011).*

There have been a handful of notable cases in which researchers have used online games to crowdsource solutions to big biological problems. The most famous of these, Foldit, relies on players (or citizen scientists, as they are known in research circles) to find novel ways of folding proteins. EteRNA, developed by scientists at Stanford and Carnegie Mellon University, does something similar with RNA molecules; and EyeWire, which was spearheaded by Princeton neuroscientist **Sebastian Seung, PhD**, has gamers map the three-dimensional structure of neurons in the retina. The results can be impressive: In a paper published last year in the journal *Nature*, Seung and his co-authors drew on the work of thousands of EyeWire players to help explain how eyes detect motion.

Still, games that are used to advance biomedical research remain rare enough to garner attention just for their sheer novelty. But that may be about to change.

Last December, the National Institutes of Health (NIH)sponsored a two-day workshop at the National Cancer Institute's (NCI) Shady Grove facility in Rockville, Maryland. On the agenda: exploring how games could be used for biomedical research, and how the methods and technologies that game developers rely upon might also be exploited by scientists.

On the agenda: exploring how games could be used for biomedical research, and how the methods and technologies that game developers rely upon might also be exploited by scientists.

The benefits could be substantial. The avalanche of data available to researchers is fueling demand for new and better tools to analyze and understand it—tools that use methods such as crowdsourcing and data visualization to discover patterns and solutions that might otherwise go unnoticed. As it turns out, game developers have been refining such methods for years.

Bringing Game Thinking to Biomedicine

"Game developers are computer scientists who just happened to go into a slightly different field," says **David Miller, PhD**, AAAS Science and Technology Policy Fellow at the NCI and one of two lead organizers of the December event. His fellow organizer, **Jennifer Couch, PhD**, who leads the Structural Biology and Molecular Applications Branch at the NCI, adds that biomedical researchers have watched with considerable envy as game developers have succeeded in visualizing and manipulating complex systems for

entertainment purposes, getting large numbers of people to work together online.

Riedel-Kruse, meanwhile, points out that game developers are also highly adept at motivating players, designing incentives to keep them engaged, and building interfaces that help them learn how to play as they progress through the game itself. Both science and games involve problem-solving, he says, but game developers are the ones who have figured out how to encourage voluntary participation by making their products fun, even if they may also be frustrating. That skill could prove useful not only for building online research-oriented games capable of marshaling thousands of citizen scientists, but also for developing digital tools to make the more tedious tasks involved in biomedical research—like tracking the positions of individual cells in a sample—more bearable for professional scientists toiling away in their labs.

In EteRNA, developed by scientists at Stanford and Carnegie Mellon, players design complex new RNA molecules. The best designs are then synthesized in the lab.

The workshop, which was sponsored by the NIH's Big Data to Knowledge (BD2K) initiative—a program that aims, among other things, to develop the methods and tools necessary to analyze biomedical Big Data—and formally titled "BD2K Think Tank: Game Developers and Biomedical Researchers," was therefore intended not only to generate ideas for new games that could help solve specific problems, but also to explore how scientists might benefit from what game developers sometimes call game thinking. "The way that game developers think about problems is different from the way in which bioinformaticians think about problems," Couch says. "How can we bring some of that thinking into biomedical research?"

Making that happen will require a good deal of effort on both sides. The workshop participants—nine biomedical researchers and ten game developers—were all chosen in part for their open-mindedness, but that didn't mean they knew much about their counterparts' work or methods. Indeed, just getting up to speed on the basic science involved in biomedical research is going to be a challenge for game developers, says **Ben Sawyer**, a leading figure in the realm of serious games who co-chaired the event and recruited the developers—though he adds that they're up to the task, and even look forward to it. (Sawyer co-founded Games for Health, a grassroots network funded by the Robert Wood Johnson Foundation that supports the development of health games and technologies.) Researchers, meanwhile, are going to have to get their minds around an entirely different way of approaching problems. Whereas scientists tend to begin by looking at big questions and try to design experiments that will answer them, says Couch, game developers work more iteratively, jumping right into a project and allowing the solution to reveal itself—often in a form that they might not have envisioned. "We're just worlds apart," says **Markus Covert, PhD**, an associate professor of bioengineering at Stanford who rounded up the researcher contingent and acted as co-chair.

Gamifying Research

In an effort to bring those worlds a bit closer together at the workshop, Covert and Sawyer engaged in a sample dialogue about Covert's work, which involves computational models of cellular activity and plenty of live-cell imaging. Their back-and-forth led first to a lively discussion among the assembled game developers about how Covert's research could be "gamified," followed by a round of speed-dating sessions that teamed each researcher up with several game developers to brainstorm how games might be used to help solve problems involving large data sets. The topics ranged from genomics to organic chemistry, and the sessions touched upon everything from the objectives of cancer researchers to the potentially useful characteristics of games like Pokémon and SimCity.

In some cases, game developers were able to quickly see the potential for turning particular research problems into games. And participants on both sides walked away having made connections that could lead to future collaborations. Riedel-Kruse, for example, met a number of developers whom he believes could potentially help refine his biotic games; while **Nick Fortugno**, a prominent game developer and entrepreneur who teaches game design at The New School's Parsons School of Design, is already in discussions with several different researchers who attended the meeting. But perhaps just as important, the various parties involved also left with a better understanding of the opportunities—and challenges—that lie ahead.

Both science and games involve problem-solving, he says, but game developers are the ones who have figured out how to encourage voluntary participation by making their products fun, even if they may also be frustrating.

Sawyer, for example, points out that research-oriented games pose a unique problem for game developers: whereas the latter typically know all of the rules of a game before they design it, they won't have that luxury when constructing games meant to facilitate scientific discovery—games whose very purpose will be to help lift the veil on the unknown. Yet building games in the absence of all the rules not only flies in the face of traditional game design; it could also result in games that upend traditional norms of gameplay. "Imagine if you only knew 30 percent of the rules of chess," says Sawyer, "and I arranged the board in some weird way that was still valid, and I said, 'How did we get there?'"

Fortugno explains that this kind of uncertainty about basic ground rules will make it harder for game designers to evaluate the solutions that players come up with, and therefore to design incentives that will keep them engaged. And the fun won't stop there: Once they've succeeded in establishing evaluative criteria based on the limited information that researchers can provide, the developers will still have to gamify those criteria in ways that will not only motivate players to win, but also produce scientifically relevant results. "This," says Fortugno, "is an odd way to make games."

On the other hand, researchers and developers alike see great potential in open-ended, discovery-oriented games. Among other things, such games could be used to test researchers' hypotheses and models by having large numbers of people run amok in gamified versions of them, identifying aberrations and inconsistencies as only game-obsessed players can. "That's the process of game design," says Couch. "You build a game and put it out there, and players find all the glitches and exploits very quickly, figuring out where it's broken." Similarly, research games could be extremely useful tools for attacking nebulous problems involving large piles of data. "We as humans are very, very good at seeing anomalies given the right type of data, presented in the right way," Miller says. And Sawyer contends that that the unusual challenge posed by research games could spur developers to discover new ways of designing systems to deal with unknown rules and data structures, resulting in games—and outcomes—that go beyond anything achieved thus far.

Where this may lead, only time will tell; and there will be plenty of practical hurdles to overcome, like finding ways of connecting scientists and game developers who inhabit completely different professional networks, and figuring out how to fund their collaborations. But Couch and Miller are already considering ways of continuing the conversation between researchers and developers, perhaps through boot camps or short courses. And in a telling sign, the BD2K initiative recently announced a new funding opportunity to support the development of "new or significantly adapted interactive digital media that engages the public, experts or non-experts, in performing some aspect of biomedical research via crowdsourcing."

Let the games begin. □

Game developers who attended the NCI think tank pondered whether popular, open-ended games such as SimCity (pictured) and Minecraft could serve as models for biomedical research games.

...Research games could be extremely useful tools for attacking nebulous problems involving large piles of data. "We as humans are very, very good at seeing anomalies given the right type of data, presented in the right way," Miller says.

VIRAL PHYSICS LESSONS:
Simulations Offer Novel Insights

By Kristin Sainani, PhD

Zooming in on a virus reveals a physical marvel. It can stuff a genome into a confined space (a protein casing called a capsid). It can eject its genome rapidly and fluidly into a cell. And it can coat itself in a piece of host cell membrane (a lipid envelope) to avoid detection by the host's immune system. To better understand how viruses perform these parlor tricks, some scientists are turning to physics-based computer simulations. These methods are revealing some unexpected vulnerabilities in viral design that could lead to novel ways of crippling these creatures as well as ways to engineer viruses for therapeutic purposes.

Spring-Loaded DNA

To cram their genetic material inside a capsid, DNA viruses employ some of the most powerful molecular motors in nature. It takes a lot of energy to bend the stiff double helix and overcome repulsive forces between strands. Once squished inside the capsid, the genome waits ready to spring out at the moment of infection, a process called DNA ejection. "We suspected that there is a one-to-one relationship between the way the DNA is packaged by the motors and the dynamics of ejection," says **Murugappan Muthukumar, PhD**, professor of polymer science and engineering at the University of Massachusetts, Amherst.

> To better understand how viruses perform [their] parlor tricks, some scientists are turning to physics-based computer simulations.

His team explored the connection between packing and ejection in bacteriophage in a 2013 paper in *Biophysical Journal*.

Muthukumar's team modeled DNA as

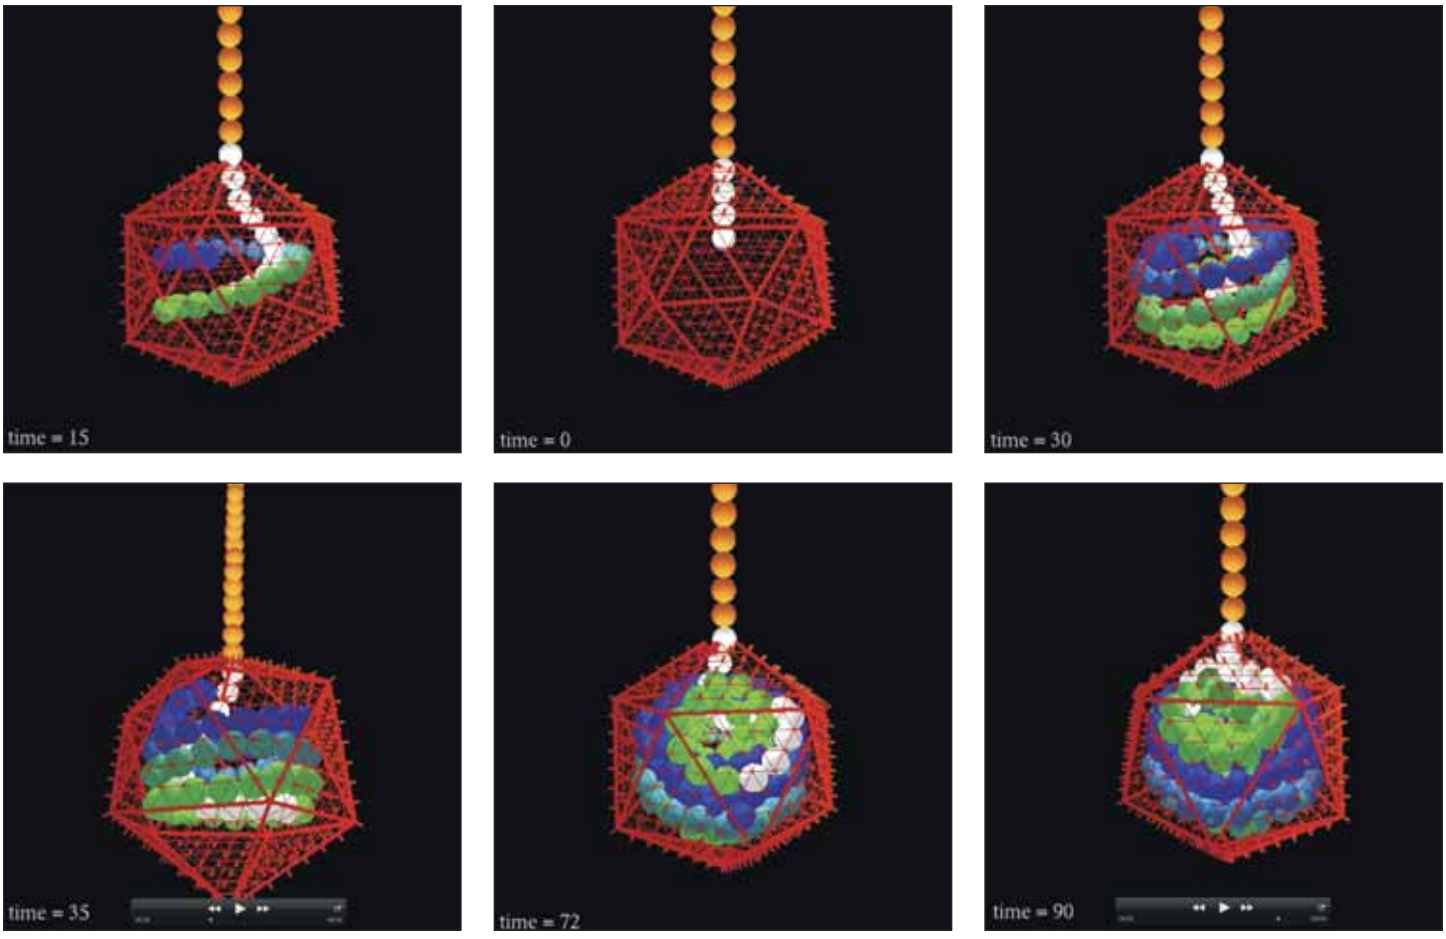

These screenshots show the progress of a DNA viral packing simulation from initial entry of the DNA to near completion. In repeated simulations, the researchers found that DNA organization after packaging varied and also depended on motor force. Courtesy of: Murugappan Muthukumar.

a wormlike chain of charged beads on a string. In their simulation, the motor stuffs the rope-like molecule through an opening in the capsid; after packing is complete, the DNA is allowed to spontaneously eject from the same portal. As experimentalists had previously observed, they found that the motor frequently stalls during packing. Experimentalists have blamed the pauses on motor choppiness, but Muthukumar's team showed that the DNA chain sometimes bunches up, and takes time to relax and make room.

They also showed that packing occurs slightly differently every time and is more variable and disorganized at higher motor speeds. The more ill-ordered the packing, the more inconsistent the ejection kinetics. "We watched each DNA monomer as it was coming out. We could see whether it was wrapping around, going in the wrong direction, or going in an orbital-like movement." They found that if DNA approached the exit portal at the wrong angle, it would become jammed, causing ejection to pause as the molecule straightened out. A 2014 paper in *PNAS* confirmed their predictions with experimental data.

These observations could be exploited therapeutically. For example, therapies could be developed to speed up the molecular motor enough to make packing more disorderly, causing the DNA to become stuck in the capsid. Or treatments might mess with the DNA inside the capsid, preventing it from achieving an effective exit path, Muthukumar says.

...Therapies could be developed to speed up the molecular motor enough to make packing more disorderly, causing the DNA to become stuck in the capsid. Or treatments might mess with the DNA inside the capsid, preventing it from achieving an effective exit path, Muthukumar says.

Self-Compacting Genomes

Unlike DNA viruses with their powerful molecular motors, many single-stranded RNA viruses pack their genomes with no help. Their RNA strands fold spontaneously into space-saving structures, and the protein capsid takes shape around them. **Luca Tubiana, PhD**, postdoctoral fellow at the University of Vienna, wondered how RNA viruses manage that task. "Viral RNA are more compact than completely random RNA," he observes. So, he hypothesized, perhaps their compactness is the result of evolution.

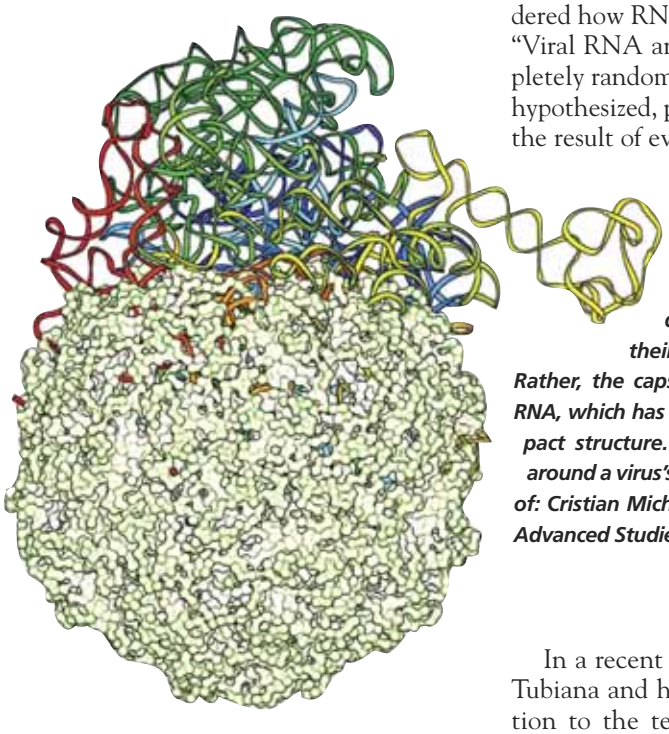

Single-stranded RNA viruses don't require a motor to pack their genomes into the capsid. Rather, the capsid self-assembles around the RNA, which has evolved to have an extra-compact structure. Here, capsid proteins gather around a virus's single-stranded RNA. Courtesy of: Cristian Micheletti, International School for Advanced Studies.

In a recent paper in *Biophysical Journal*, Tubiana and his colleagues put that question to the test. The team mutated the genomes of 128 different single-stranded RNA viruses *in silico* using only synonymous mutations, which do not change the amino acids encoded. Then they used software to predict how the RNA would fold. Strikingly, mutations in just five percent of the genome were sufficient to wipe out genome compactness. This was true even if they preserved codon bias (the fact that some codons occur more frequently than others) and restricted mutations to protein-coding regions. "This strongly indicates that this compactness is evolutionarily selected for," Tubiana says. These viral RNAs don't just code for specific proteins but also for physical shape, he concludes.

This knowledge could be useful in the design of antiviral drugs. "In principle, one may be able to destroy this physical compactness and therefore hinder the reproduction of the virus," Tubiana says. Plus, scientists are hoping to exploit viruses for therapeutics, such as engineering anti-cancer viruses. Understanding how capsid assembly works naturally will make this task easier.

The Physics of a Virus's Shield

Many of the world's deadliest viruses—including Ebola, HIV, and flu—are encased in a lipid envelope that shields them from the host immune system. Newly formed viruses steal pieces of the host cell membrane, into which they insert the viral spike proteins that will be used to latch onto new host cells. In a recent paper in *Structure*, scientists reported the first microsecond-timescale simulation of an influenza A virion's envelope. Their goal: To gain insight into the flu virus's biophysical behavior, says lead author **Tyler Reddy, PhD**, a postdoctoral fellow at the University of Oxford.

They based their models on abundant experimental data, including viral protein structure as well as the lipid composition of the envelope. "The details, the shapes of the models, come from experiments. But these experiments are largely static. With computer simulation, we can animate the system," Reddy explains. "Basically, Newton's laws of motion are being used to allow the atoms to wiggle and jiggle over time."

"The details, the shapes of the models, come from experiments. But these experiments are largely static. With computer simulation, we can animate the system," Reddy explains. "Basically, Newton's laws of motion are being used to allow the atoms to wiggle and jiggle over time."

His team first modeled the envelope as a lipid ball in a water droplet. Once the ball (vesicle) relaxed into equilibrium, they inserted viral spike proteins into the model and capped some lipids with sugars to fashion the glycolipids that account for 12 percent of the envelope. Though the model was coarse-grained to reduce computational demands, their five-microsecond simulations still took a year to run on a high-performance supercomputer.

The simulations revealed that the surface glycolipids slow down the movement of both spike proteins and lipids. "That makes sense because they're basically these physical obstructions on the surface of the virus," Reddy says. Low envelope mobility may help explain how flu viruses can survive in water for up to three years, he says. Reddy's team also showed that the spike proteins don't clump in the presence of glycolipids, which likely facilitates host cell binding but may also make the virus vulnerable to host antibodies.

Reddy's team is working on a bigger simulation that incorporates both the flu virus and the host cell membrane to see how they interact. Eventually, they hope to use their model to probe how flu virions respond to different drugs.

These computer simulations reveal new insights into how viruses reproduce and spread. Beyond the implications for medicine, viruses' slick anatomies also evoke wonderment. Muthukumar says: "Viruses are very beautiful objects." □

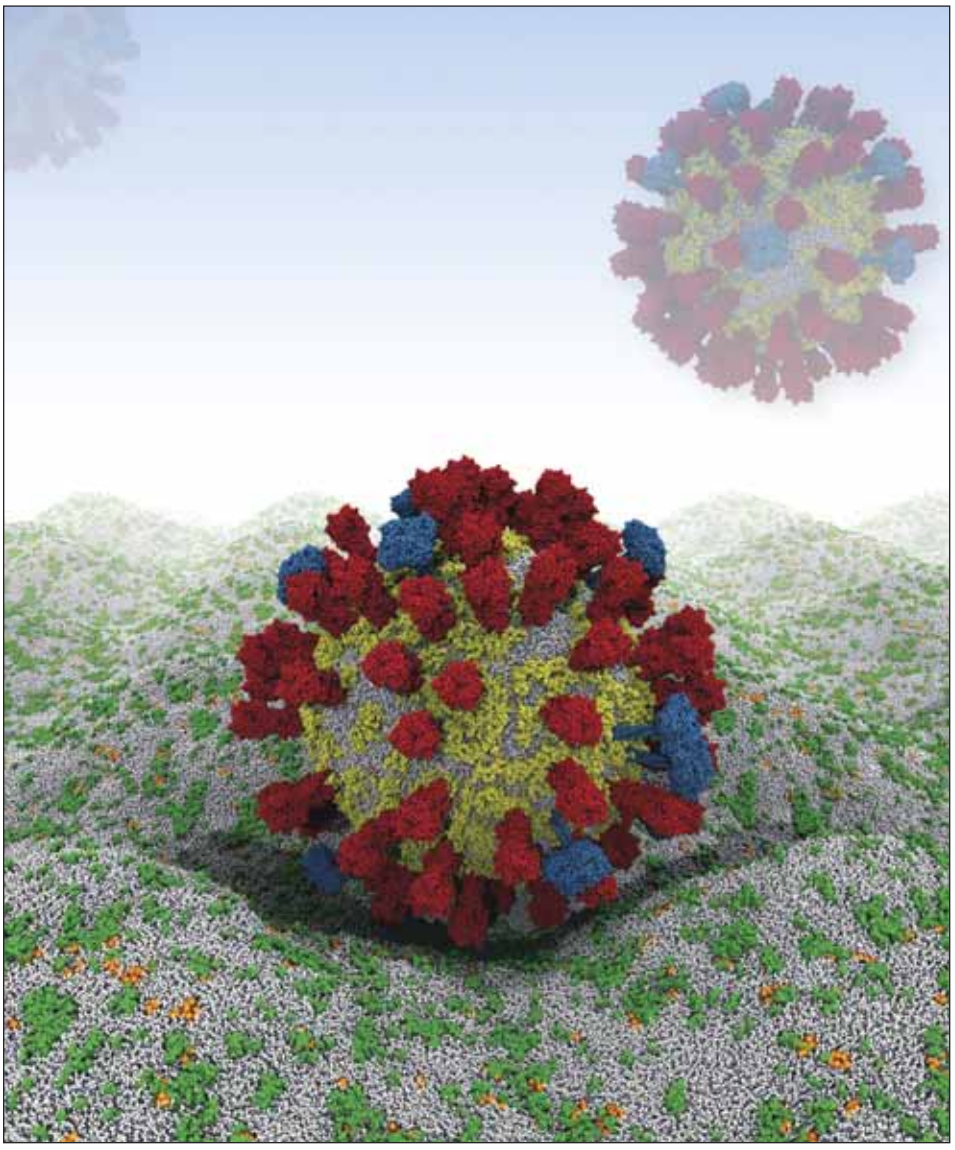

In the future, scientists from the University of Oxford plan to build on their flu envelope model; for example, they may simulate its interaction with a host cell membrane, as pictured in this cartoon. Courtesy of: Heidi Koldsoe, University of Oxford.

CITIZEN SCIENCE:
Getting Cheap, Reliable Help from Lay Workers

By Esther Landhuis

In pre-Internet days, people sought expert advice for their purchasing decisions—consulting product ratings in magazines such as *Consumer Reports* and reading newspaper reviews of local dining spots. Now, many rely instead on feedback from ordinary folks who post to websites like Amazon or Yelp. The soaring popularity of such online reviews signals the value of the crowd: Even if any single lay opinion might seem dubious, the wisdom of the group offers a powerful source of information.

Biomedical researchers have watched these developments eagerly. "Those of us on the interface between technology and biology were thinking, 'Hey, how can we apply crowdsourcing to problems and challenges that we care about?'" recalls bioinformatician **Andrew Su, PhD**, of Scripps Research Institute in La Jolla, California.

Data Deluge

In recent years, the potential value of crowdsourcing has grown as biomedical researchers drown in data. Advances in online technology and DNA sequencing have plunged biomedicine into a new era in which genome-scale analyses churn out data faster than anyone can make sense of it.

So, for the last decade, Su's research group has reached out to crowds to organize biomedical information. One of their first projects was Gene Wiki, a collection of 10,000-plus pages with information on human genes and proteins. The portal is hosted on Wikipedia, the free online encyclopedia that made a splash in 2001 and is now the world's sixth-most-visited website with 34.5 million articles written by more than 53 million people worldwide. "None of their content is created by paid individuals. It's all on the backs of volunteers," Su says. "It speaks to the power of crowdsourcing."

The idea for Gene Wiki emerged as genome-scale experiments became more powerful and feasible. It's not uncommon now for a single experiment to produce a list of some 500 genes expressed at different levels in cancer versus normal cells. While a researcher might know something about one or two of those genes, Su notes, "I need to quickly get up to speed on the other 498 that I'm not familiar with—to understand if they're relevant to my system or worth further study."

Resources like Gene Wiki require tons of biocuration—the process of combing through biomedical literature and putting its content into structured databases that can be queried for statistics and trends. The National Institutes of Health spends millions of dollars each year hiring professional scientists to do biocuration. "We hope to make that process more efficient by engaging crowds," Su says. "The more we can get our crowd to do, the more professionals can focus on really hard problems."

In the training phase of the Mechanical Turk project, lay workers were given feedback on how well they did on a task, such as identifying disease names. Reprinted from BM Good, et al., Microtask Crowdsourcing for Disease Mention Annotation in Pubmed Abstracts.* Biocomputing 2015*: pp. 282-293.

Su's eventual goal is to build a Network of BioThings. This system would organize the torrents of data that currently flood the PubMed database at a rate of one or two new articles per minute. "Keeping up with the literature is incredibly hard," Su says. Rather than spending weeks scouring abstracts, it would help if researchers could glean their useful tidbits by querying a knowledge base. Building it would involve surveying publications for key "bio things"—genes, proteins, mutations, diseases and drugs—and documenting relationships between them.

Building and Debugging Databases

As a first step, Su and colleagues tested if they could crowdsource this sort of biocuration to lay people. Using Mechanical Turk—a web platform for harnessing human intelligence to do things computers can't do well—they asked workers to highlight disease-related terms in 593 PubMed abstracts.

This job had previously been done by 12 professional biocurators working part-time for a good part of the summer, an effort that Su estimates cost tens of thousands of dollars. With crowdsourcing, 145 lay workers completed the work in nine days. Each document was scanned by 15 novices who earned six cents per abstract. The upshot: Six

"The more we can get our crowd to do, the more professionals can focus on really hard problems," Su says.

novices in aggregate did as well as, if not better than, one PhD biocurator, and at a fraction of the cost ($631 total including time for training). Su reported these results at the Pacific Symposium on Biocomputing (PSB) held January 4 – 8 in Hawaii. And Su has recently shown that the same work could be done just as reliably—and free of charge to the researchers—on the Mark2Cure.org site, which allows interested people to volunteer their time to contribute to research. In May, Mark2Cure launched a biocuration campaign aimed at aiding rare disease research.

In addition to building databases, crowdsourcing can help fix them. **Mark Musen, PhD**, and **Jonathan Mortensen** of Stanford University sought crowd help to find errors in SNOMED CT, a set of clinical terms and concepts becoming more critical as hospitals switch to electronic medical records. For example, typing "Tylenol" into a SNOMED system could warn a physician to avoid prescribing this medication to a patient who is allergic to its key ingredient, acetominophen, which SNOMED knows is the main ingredient of Tylenol.

Asked to verify relationships and find mistakes in a subset of SNOMED CT terms, lay workers performed "on par with experts" and cost one-fourth as much, Mortensen says. He was the lead author on a November 13 paper reporting these findings in the *Journal of the American Medical Informatics Association*.

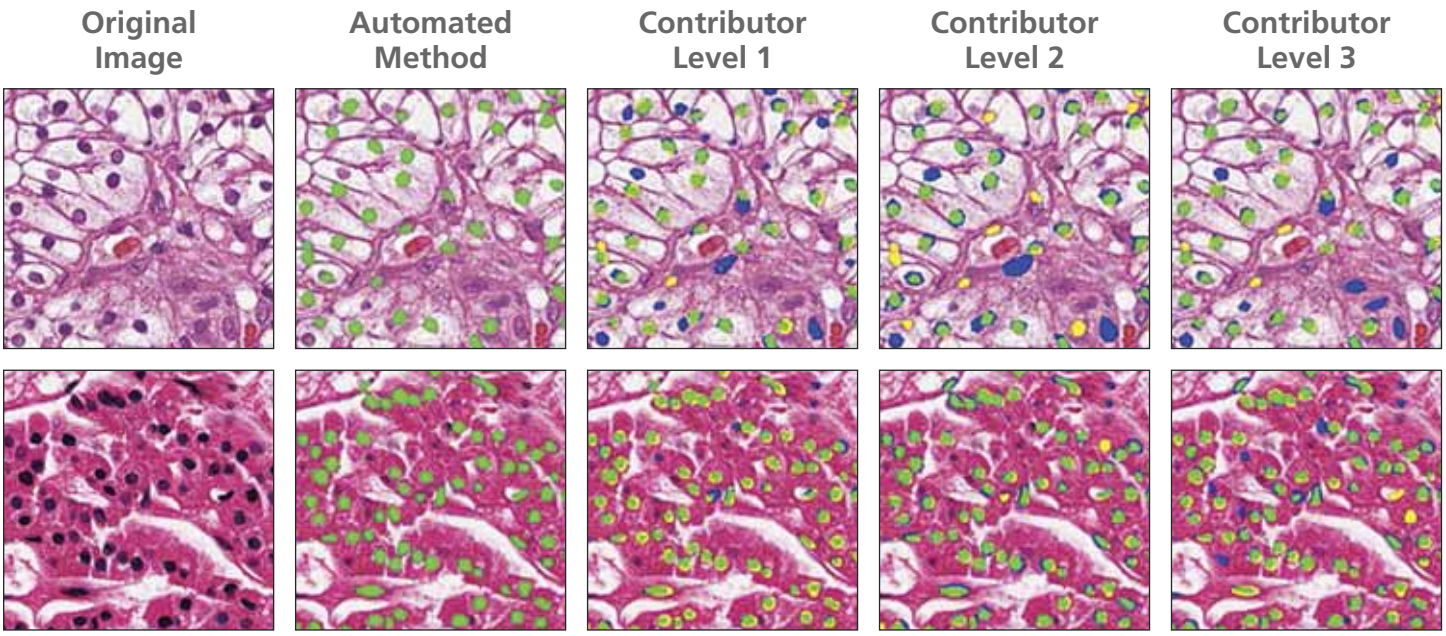

Lay workers assigned the task of determining nuclear boundaries of potentially cancerous cells did a better job than state-of-the-art automated methods. These slides show two examples (top and bottom) of nuclear segmentation using an automated method and three increasing contributor skill levels. Green region indicates a true positive region, yellow region indicates a false negative region and blue region indicates a false positive region. Reprinted from H. Irshad, et al., Crowdsourcing image annotation for nucleus detection and segmentation in computational pathology: Evaluating experts, automated methods, and the crowd,* Biocomputing 2015*, Proceedings of the Pacific Symposium, pp. 294-305.

Many Eyes Make Light Work

Crowdsourcing has also proven useful for annotating images—a huge need in cancer research. Despite tremendous advances in molecular biology that allow researchers to probe thousands of genes and proteins within individual cells, "the single most useful tool for diagnosing cancer is a microscope. It's the convergence of all this complex molecular data," says **Andrew Beck, MD, PhD**, a molecular pathologist at Beth Israel Deaconess Medical Center in Boston.

Training computers to correlate molecular and microscopy data could help physicians tell if a tumor is benign or malignant, or predict how it might respond to treatment, Beck says. The challenge is getting enough annotated images to build the algorithms.

For a study he reported at the January 2015 PSB, Beck and colleagues showed lay workers a set of renal cell carcinoma images and asked them to identify and delineate the boundaries of nuclei, which contain the cell's DNA. The size and shape of the nucleus, as well as how dark or light it appears under a microscope, can help researchers distinguish cancer cells from normal tissue.

On the first task —identifying nuclei— automated approaches did about as well as the crowd. However, for determining nuclear boundaries, human eyes did considerably better than state-of-the-art methods, Beck says.

Though this study only focused on two small tasks, Beck says the approach could be extended "to something as complex as making diagnoses." However, he notes, computers won't replace human expertise. When a computational method adds value, more people want to use it—which then creates more complicated results to be interpreted. "Ironically, the better the machines we have, the greater the need for human experts," Beck says. □

On the first task — identifying nuclei — automated approaches did about as well as the crowd. However, for determining nuclear boundaries, human eyes did considerably better than state-of-the-art methods, Beck says.

CROWDSOURCING TO EXPERTS: The DREAM Challenges

By Esther Landhuis

At a 1906 county fair in England, some 800 villagers tried to estimate the weight of an ox. None of the contestants hit the mark, but a closer look at their guess cards led to a stunning discovery. Stacking the estimates from lowest to highest, the middlemost value came within 0.8 percent of the ox's butchered weight—closer than individual guesses submitted by cattle experts. Published in 1907, these findings on the statistical concept of median were among the first to demonstrate the wisdom of crowds.

A century later, leaders of the genomic revolution are summoning "crowds" to tackle some of the toughest problems in modern medicine. These aren't crowds of ordinary townsfolk—or even biologists, necessarily. Many train in fields such as computer science, engineering or statistics and spend far more time staring at numbers and graphs than scrutinizing cells under a microscope. They're part of a collaborative initiative called DREAM (Dialogue on Reverse Engineering Assessment and Methods). Since 2007, the group has organized more than 30 open science competitions drawing diverse experts to complex biomedical questions. Wooed by prize money and opportunities to publish their approaches in top journals, researchers around the globe have developed computational models for a variety of translational medicine challenges, including predicting drug responses and disease outcomes.

DREAMing of Better Solutions

For systems biology, the crowdsourcing concept emerged as scientists were faced with organizing huge piles of data coming out of DNA microarray experiments. Microarrays measure the expression of thousands of genes at once, comparing their levels in groups of cells under normal versus disease conditions, for instance. But massive lists of differentially expressed genes by themselves aren't that useful. Researchers want to understand how the genes are connected, such as whether they encode proteins that interact or regulate other genes, says computational biologist **Gustavo Stolovitzky, PhD**, of IBM Research and the

…Leaders of the genomic revolution are summoning "crowds" to tackle some of the toughest problems in modern medicine. These aren't crowds of ordinary townsfolk—or even biologists, necessarily. Many train in fields such as computer science, engineering or statistics and spend far more time staring at numbers and graphs than scrutinizing cells under a microscope.

Icahn School of Medicine at Mount Sinai in New York, one of DREAM's founders.

Computational scientists have assembled networks using algorithms that reverse-engineer or infer gene relationships from data. However, some worry that validating these approaches relies too much on cherry-picking. By focusing on connections that seem consistent with prior publications, "you're selecting what works for you but forgetting the ones that might not be working," Stolovitzky says.

DREAM originated as a way to evaluate these network inference algorithms. Open competitions allow participants to see which schemes work and which don't. Soon, the group realized DREAM challenges could do more than assess methods—they could accelerate research. By focusing a community of experts on a specific problem for a limited time, work that might take 10 years in a single lab could be done by the crowd in several months. Reliability also got a boost. "When we aggregate the solutions from all participants, the resulting solution is often better than the best," Stolovitzky says.

By focusing a community of experts on a specific problem for a limited time, work that might take 10 years in a single lab could be done by the crowd in several months. Reliability also got a boost. "When we aggregate the solutions from all participants, the resulting solution is often better than the best," Stolovitzky says.

Challenge: Predict Cancer Drug Responses

Several papers published last year in *Nature Biotechnology* highlight DREAM challenges aimed at developing rational approaches to predict how cancer patients respond to treatments. These days, choosing drugs involves a fair amount of guesswork, unless the patient happens to have a gene mutation known to drive that particular cancer. In one challenge, the DREAM coordinators gave teams genomic, epigenomic and proteomic profiles for 35 breast cancer cell lines as well as information on how the cells respond to treatment with a group of drugs. The teams were then asked to predict how well a different set of 18 cell lines would respond to those drugs, given only their genomic, epigenomic and proteomic profiles.

The 44 algorithms submitted by the research community suggest it is possible to develop rational approaches for predicting drug responses. However, their predictions are "not yet as good as we would like," Stolovitzky says. Asked to rank cell lines from most to least sensitive for each drug tested, the top model ordered 60 percent of cell line pairs correctly. By comparison, "a monkey doing this task would be right half the time," notes Stolovitzky.

In a related DREAM challenge, participants devised algorithms to rank 91 pairs of compounds on how strongly they enhance or sabotage each other's effects—otherwise known as synergism and antagonism. This challenge proved harder—only 3 of 31 submissions performed better than random guesses. However, the top methods were based on different hypotheses about how synergism and antagonism work—and combining them produced better results and provided insights into how drug interactions might work.

Thus, "while the results are not immediately applicable to the clinic, they begin to establish the rules and types of data needed to predict accurately the correct drug regimen," says **Dan Gallahan, PhD**, who directs cancer biology research at the National Institutes of Health (NIH) in Bethesda, Maryland. "This is the type of research needed to make precision medicine a reality."

Challenge: Predict Neurodegenerative Disease Progression

One of the more successful DREAM initiatives offered $50,000 for the computational approach that could most accurately predict disease progression in people with amyotrophic lateral sclerosis (ALS). This neurodegenerative disorder has no effective treatment, and the disease course varies widely between individuals. Most patients die three to five years after symptoms appear, but some make it 10 years past onset.

Disease variability is a big challenge for the field, says neurologist **Merit Cudkowicz, MD, MSc**, an ALS specialist at Massachusetts General Hospital in Boston. It means clinical trials need to be large for tested compounds to show an effect. The heterogeneity also suggests different biological mechanisms could be at work in patients who decline more quickly or slowly. So "maybe there will be therapies that work in some people but not in others," Cudkowicz says.

Challenge organizers supplied competitors with three months of lab test data as well as demographics and family history for 1,822 people enrolled in ALS clinical trials. DREAM teams were then asked to predict each participant's disease progression over the subsequent nine months.

The ALS challenge drew 1,073 registrants from 64 countries. Top-performing models predicted disease outcomes better than a panel of 12 experts, and the winners "didn't know anything about ALS," Stolovitzky says.

Statisticians estimate that the best two algorithms could reduce the size of ALS clinical trials by 20 percent. For a 1,000-patient Phase 3 trial, that would save $6 million. One company—Origent Data Sciences in Vienna, Virginia—is working to incorporate new predictive analytics into future ALS trials. By estimating how a patient's symptoms would progress without the intervention, these tools are particularly useful in early trials that lack placebo arms, says Origent CEO **Mike Keymer**.

Researchers won't know the true impact of DREAM algorithms for a while. But in the meantime, the challenges have succeeded in getting cross-disciplinary researchers out of their silos and working together. □

My phone buzzes.

It's Mood Matters, a mood-tracking app developed by the startup Ginger.io. "We notice you haven't logged any recent physical activity," it alerts me, linking to an article about the connection between depression and exercise. I glance at the band on my wrist, a Fitbit fitness tracker that's unrelated to the app, and see that I've only walked a measly 800 steps today. I scroll over to see my heart rate—at least I'm relaxed, I think. I scan my to-do list and then stand up for a quick walk around the block. Each step I take, eventually, is relayed to the cloud and stored as a bit of information in a data center with all the other steps people are taking around the world, forming a massive data set describing when and how we move.

Right now, that data set—the vast amount of information already collected by mobile health devices—is mostly looked at through the lens of very basic statistics to answer questions of curiosity. How many steps do American Fitbit owners walk, on average? Which occupations are most active? Whose heart rates spiked during the Super Bowl?

There's a vision for the future, though, that is far more complex.

In this future, devices on our wrists, in our phones, or tucked in our pockets are more than step-count monitors. They track all aspects of people's health and act as part-counselor, part-physician, part-coach, alerting us to health concerns or spurring us to make lifestyle changes. Patients could be alerted if they have signs of impending heart failure, worsening Parkinson's disease, or a low blood sugar crash—among many other things. Moreover, as devices funnel increasingly large amounts of information to the cloud, they give scientists a rich and ever-changing platform to use for research—letting them make new connections between facets of people's behavior and health that have never been linked before.

Many of the arguments in favor of the increasing use of mobile devices to monitor patients center around preventive medicine—the idea that many chronic diseases such as diabetes and heart disease can be prevented by changing people's diets or exercise patterns, saving healthcare systems vast amounts of money.

"It's easy to give a pill, it's straightforward to do a procedure, but to change patients' behavior is the holy grail in medicine," says **Alan Yeung, MD**, La Ka Shing Professor of Medicine at Stanford School of Medicine. "Phones, together with wearable devices, can for the first time provide some objective evidence of behavior." And once researchers understand what influences behavior they can set to work changing it.

"The ultimate goal is to improve health outcomes for people," says **Ray Browning, PhD**, associate professor of health and exercise science at Colorado State University. "We have a lot of chronic disease in this country that's preventable with changes to behavior."

Today, mobile fitness devices are exploding in popularity, but we're only a small way toward that vision coming true. "We're at the very beginning of mobile devices starting to impact how medicine is being practiced," says **Eric Topol, MD**, a Scripps Research Institute cardiologist who has written extensively about the technological future of medicine. He thinks the slope toward full adoption likely has a hockey stick shape. "We're starting to get closer to that rapid rise."

Interviews with a handful of researchers and companies who are pushing the field forward suggest that, while many of them are making lofty promises for the future, challenges remain: showing the clinical utility of devices and apps through not just anecdotes but well-designed clinical trials; getting both doctors and patients alike to buy into the use of the devices; and developing new computational methods to parse the steady stream of data from the mobile device fire hose.

Get Moving

In the decades-off vision of wearable health trackers, devices on our wrists or in our pockets can collect all sorts of data about our breathing and our eating and markers in our blood. But for now, the majority of wearable medical devices are fitness trackers that rely most heavily on one piece of data: our movement.

"The lower-hanging fruit in the field right now is physical activity data because it's so ubiquitous," says **Ida Sim, MD**, a co-director of the Biomedical Informatics Institute at the University of California, San Francisco, as well as an investigator with Mobile Sensor Data-to-Knowledge (MD2K), an NIH-funded Big Data to Knowledge (BD2K) Center of Excellence. "The thing being measured is a physical quantity and it's pretty easy to represent and to calibrate between devices."

Devices like the Fitbit, or numerous phone apps, rely on accelerometers to tell users how much they've walked each day. Tiny crystal structures embedded in these devices sense movement—as your arm swings or your body moves up and down in a tell-tale walking pattern—by detecting changes to the direction they're pulled by gravity. Then, they transmit a voltage relaying this information. It's a relatively simple and cheap technology these days—the accelerometer in the latest iPhone has an estimated cost of 65 cents and the gyroscope (a similar technology that detects the tilt of the phone) costs less than two dollars.

But—at least until now—that step count

Wearing Your Health *on Your Sleeve:*

How big data from mobile apps and sensors may revolutionize healthcare

By Sarah C.P. Williams

has rarely been linked to real health advice beyond the idea that more activity can help you lose weight and lower your overall odds of a plethora of obesity-related chronic disease. Now, though, that's starting to change. Researchers are now seeking meaningful associations in the data deluge from early adopters, with the ultimate goal of showing that their app or device of choice can make people healthier or less likely to develop conditions like heart disease.

The Fitbit dashboard shows users a variety of data about their daily activities. The data gathered could help researchers discover whether this kind of tracking can actually change users' behavior.

Yeung, with colleagues at Stanford, has developed MyHeart Counts, one of the apps that Apple is touting as the future of its smartphone and smartwatch-based technologies. When you download MyHeart Counts—as roughly 26,000 people did in the 24 hours after Apple's latest press conference on its new open-source ResearchKit—you're asked a series of questions on your lifestyle and family history of heart disease. Then, using the iPhone's movement sensors, the app records your every motion for seven days.

"Using the gyroscope and GPS on the phone, you can easily tell whether someone is just sitting around or whether they're being active," says Yeung. If someone has time, they can also take a six-minute walk test—their phone records how far they are able to walk in that time period. It's a classic test that's been used by cardiologists for years, but usually relies on a stopwatch and marks along the edge of a long hospital corridor. With the phone app, users can perform the test anywhere, and at the end, MyHeart Counts provides each user with a calculated "Heart Age." "A person signing up might be fifty years old, but we might calculate that their heart health is more like that of a sixty year old," Yeung says. Again, this computation is not new—it is based on existing data published by the American Heart Association—but the app makes it more accessible and understandable to the average person. Yeung hopes it will motivate people to set a goal of lowering their heart age.

Like many other devices and apps, though, the challenge with MyHeart Counts is proving any sort of clinical utility. Generally, to show that patients benefit from an intervention (whether it's a drug, a counseling session, or a device), researchers have to show that people who use it have less disease—or longer lives—than those who don't.

For now, Yeung and his colleague are tracking whether those who continuously use the app—and get reminders to exercise—see a decrease in their computed heart age. Eventually, they'd like to test out the app in patients who are at higher risk of heart attacks, for example those who have diabetes or have had heart bypass operations. Could an app tracking their behavior signal which of these patients need doctors' visits or new medication strategies?

"If we evolve the app to be more clinically relevant, we'd like to give it to every heart patient at Stanford," Yeung says.

The question remains whether reminders on an app or device are sufficient to motivate behavioral change, says Browning, who collaborates with Stanford's Mobilize Center on a project that will apply new analytical techniques to movement data from devices and apps.

"I've never heard of a person who says, 'The reason I don't have a more active lifestyle is because I don't have a wearable activity monitor to tell me how active I am,'" Browning points out.

He says that it will take massive public health campaigns—ones he likens to anti-smoking campaigns—to get activity on people's radar. Until then, the majority of the population that already doesn't spend

Continued on page 18

The MyHeart Counts app uses the Apple ResearchKit (see sidebar) to collect data on users' cardiovascular health and send them to Stanford researchers who hope to learn how to motivate users to become more active.

An App a Day KEEPS the Doctor Away

This spring, consumer electronics giant Apple went from being an intermediary in the mobile health market—their iPhones could track steps or host third-party fitness apps—to being a major stakeholder. Apple's jump into the big time could mean good things for those who want to see mobile health go mainstream.

"Personally, I'm excited to see the Apples and Googles and Samsungs of the world take on preventive healthcare," says Ray Browning of Colorado State University. "You're talking about a lot of horsepower all of a sudden being thrown at these problems. And these companies are historically highly successful at changing behaviors."

In March, Apple unveiled ResearchKit, which aims to transform the way data for clinical trials is collected. ResearchKit launched with five apps—including MyHeart Counts which tracks cardiovascular disease (see main story). Other apps log patients' symptoms of diabetes, Parkinson's disease, asthma, and breast cancer. But the power behind ResearchKit is that the data isn't just for patients' own curiosity—it's linked to research programs that want to use the data to answer questions about diseases. The Wall Street Journal reported that in the first month of availability, more than 60,000 patients signed up for the apps—essentially volunteering themselves as clinical subjects.

Then, in April, Apple announced a second major collaboration: a partnership between Apple, IBM's new Watson Health Cloud, Johnson & Johnson, and Medtronic. Together, the companies are launching an effort to tailor data storage and data analytics to clinicians. The data they try to harness will be drawn not only from ResearchKit, but from Apple's earlier HealthKit, the basic suite of fitness apps that comes pre-installed on iPhones.

Already, a survey by Reuters news agency found that more than half of the top US hospitals have rolled out pilot programs using HealthKit. With patients' consent, doctors can view data collected by their smartphones and add it to electronic medical records. Google and Samsung have also launched collaborations with hospitals and medical record providers to pair their software more closely with clinicians. Like all mobile health efforts, however, it remains to be seen whether buy-in from these major tech companies can not only streamline the way data is collected, but change patient behaviors. □

"Personally, I'm excited to see the Apples and Googles and Samsungs of the world take on preventive healthcare," says Browning. "You're talking about a lot of horsepower all of a sudden being thrown at these problems. And these companies are historically highly successful at changing behaviors."

much time exercising isn't about to pick up a fitness tracker and start working out, he says.

Inside Your Head

If tracking movement with the aim of preventing obesity, diabetes, and heart failure is the low-hanging fruit in mobile medicine, then tracking movement with the aim of detecting downward spirals in depression is the next branch up. People who are depressed are less likely than usual to call and text friends, more likely to stay home, and less likely to exercise. And these are all things that can be detected by a smartphone's call logs, text message records, and gyroscope using an app such as Mood Matters from Ginger.io.

"Of course everybody deals with depression slightly differently," says **Joe Grimberg**, head of marketing at Ginger.io. "But clinically we know that social isolation and physical lethargy are markers of depression."

By giving clinicians access to day-to-day trends in patients' activity levels, the app is designed to detect when caregivers should intervene to help patients who have depression, as well as guide psychiatrists' conversations with those patients, Grimberg says.

They've partnered with psychiatrists at the University of California, Duke, and other medical research centers to test out what happens when clinicians can see the daily behavior patterns of their patients. Each institution has launched slightly different trials using the app, targeting different patient populations. Those pilot studies are ongoing. At the University of California, San Francisco, teams of psychiatrists and nurses are interacting with hundreds of depressed people who were recruited online, treating them and following their moods with no face-to-face visits. Instead, doctors and nurses receive alerts if certain behaviors are flagged—if a patient reports that they're hearing voices, or feeling suicidal for instance.

For now, the flagged behaviors that Ginger.io uses to alert a medical team that it's time to call a patient are based on previous knowledge of psychiatric disorders. Eventually, though, Ginger.io could use their data to find new "flags" that inform clinicians of how their patients are doing. So far, most of the data collected by Ginger.io—showing that the app can, in fact, gauge depression severity—has been published in the form of patent filings or presented at meetings. The Ginger.io team and their collaborators, however, are aiming to publish more recent studies on the effectiveness of interventions in peer-reviewed journals. In pilot studies, Grimberg says, they've gotten overwhelmingly positive feedback from both patients and providers.

"The great thing about the smart phone as a medical device is that it's already a part of most people's daily life," Grimberg says. "Ninety percent of smart phone users have their phone within three feet of them all day."

"The great thing about the smart phone as a medical device is that it's already a part of most people's daily life," Grimberg says. "Ninety percent of smart phone users have their phone within three feet of them all day."

And an app like Mood Matters is more

Ginger.io is teaming with various research groups to determine whether their Mood Matters app can help physicians intervene to help people with depression. Courtesy of Ginger.io.

accurate at capturing a patient's mood than a questionnaire in a psychologist's office, the company has found—not only because of its activity meter, but because of its built-in questions and journaling function, which give more frequent insight into a patient's feelings compared with reports gleaned at occasional in-person appointments.

Data from All Sources

Alex Markowetz, PhD, a computer scientist at the Universität Bonn in Germany, says that it's not just smartphones that hold the power to reveal someone's behavioral trends—it's all the computers we interact with every day.

"Theoretically, we can track any human-machine interactions," Markowetz says. "A smartphone works, but so does your email client, or World of Warcraft, or Skype, or your car or your fridge."

Markowetz's group has already developed an app called Menthal Addicted that tracks people's smartphone usage, hoping to pinpoint factors that make people more prone to becoming addicted to technology. Now, they're teaming up with clinicians to make Menthal Depressed, which, like Ginger.io, uses a person's phone to detect depression; Menthal Skilled, which requires users to complete a simple cognitive puzzle to unlock their phone and could pinpoint the earliest signs of dementia; and Menthal Dopa, which uses a phone's accelerometer and gyroscope to detect hand tremors and track the severity of Parkinson's disease over time.

All the Menthal projects, Markowetz says, are geared around the idea of getting more constant data about a patient's life to inform a clinician.

Right now, the typical clinican has two ways to get information about a patient: a questionnaire or an office visit. But these isolated data points present a tiny fragmented view of mental health, Markowetz says. "They are just poor ways to steer medication or therapy."

To develop each Menthal app, Markowetz closely analyzes how skilled doctors assess the symptoms of patients they see. He then tries to recreate that using sensors and the right computational analysis.

"A doctor might say, 'I noticed that a patient was slouching and looking downward and that made me think they were depressed'," Markowetz says. "So then I say, 'okay what sensors can we use to determine someone's posture and the direction of their gaze?'"

But finding such sensors isn't necessarily easy. For example, there's no device to measure slouching right now, although Markowetz suggests that a head-mounted device like Google Glass might work. In any event, it can take many iterations of programming and analysis to fine-tune the data signature that points toward a behavior—differentiating slouching from nodding, or eating from drinking, for instance.

And making technology applicable to healthcare requires more than the right sensors and computational methods, Markowetz adds. Sensors need to be unobtrusive and require little buy-in from the patients. If you ask someone to fill out a questionnaire on their smart phone every day 365 days a year, they're going to get tired of it quickly. Likewise, if you ask someone to wear a soap-bar-sized sensor on their forehead while they sleep, it likely won't last long. It will take small devices that require no input from users to truly make mobile medicine applicable to the whole population, Markowetz says.

Beyond Step-Counts

The technology to detect a person's steps may now be ubiquitous, but to make mobile medicine more broadly useful—beyond the prevention of obesity and depression—will take not only the most unobtrusive sensors, but sensors that can detect physiology and measures of health beyond movement.

About five years ago, **Emre Ertin, PhD**, an electrical engineer at Ohio State University who is also part of MD2K, developed AutoSense, a sensor suite that contains a one by two inch sensor that is worn on a person's chest. It is designed to track levels of stress by measuring the electrical activity of the heart as well as a person's breathing rate, temperature, and movement. As it records the data, AutoSense streams it to a cell phone.

In 2012, Ertin's research team turned to developing a stress sensor that works wirelessly without touching the skin. The device they came up with could fit in a pocket and used radio waves to sense a person's heart and lung motion. But there was a problem: The sensor was too sensitive to changes in water content. "Wireless waves don't move well through water," Ertin says. For measuring stress, this was an annoyance: The team had to figure out how to make the sensor give consistent readings even if the body's liquid levels changed. "But then we got an idea," Ertin says. "Maybe we could use the sensor to monitor lung fluid levels in congestive heart failure patients."

Almost a quarter of patients hospitalized for congestive heart failure are rehospitalized within a month, and more than half within six months; it's a number that doc-

All the Menthal projects, Markowetz says, are geared around the idea of getting more constant data about a patient's life to inform a clinician.

Right now, the typical clinican has two ways to get information about a patient: a questionnaire or an office visit. But these isolated data points present a tiny fragmented view of mental health, Markowetz says. "They are just poor ways to steer medication or therapy."

"Theoretically, we can track any human-machine interactions," Markowetz says. "A smartphone works, but so does your email client, or World of Warcraft, or Skype, or your car or your fridge."

tors are always trying to lower. But it's tough to predict which patients will have recurring problems. A sensor monitoring fluid in the lungs (a telltale sign of heart failure) could help pinpoint these patients early.

As part of MD2K, Ertin's group is collaborating with clinicians to use the wireless monitors to track patients with chronic obstructive pulmonary disease (COPD). Like heart failure, COPD is characterized by fluid in the lungs.

Some at MD2K, using sensors similar to those Ertin has developed, are working on detecting when cigarette users smoke. "There's a very particular gesture and breathing pattern that goes with smoking," Ertin says. MD2K-affiliated researchers at the University of Massachusetts developed RisQ, which uses a wristband to detect smoking behavior with 95.7 percent accuracy. Such data might provide a patient (and his or her doctor) with hints as to when they're most likely to pick up a cigarette, he says, helping tailor interventions that work.

Many challenges remain. At the 2014 International Conference on Information Processing in Sensor Networks, for example, **Santosh Kumar**, **PhD**, of the University of Memphis—a long-time collaborator of Ertin and the director of the MD2K center—presented data on attempting to use his sensors to detect when drug addicts took a dose of cocaine. But training a device to recognize cocaine use wasn't quite as easy as smoking. A person's heart rate, blood pressure, and breathing patterns change when they use cocaine. But they also change for any number of other reasons—exercise, fear, stress, and other drugs for instance. The data Kumar's team collected—922 total days (over 22,000 hours) of data from drug users—was incredibly noisy. And they had to figure out how to clean it up enough to find meaningful trends without losing the signal of the cocaine use. It took multiple iterations of modeling, statistical analysis, and data processing to get the data to this point—and in order to detect all cocaine events, the researchers still had a false positive rate of 1 per day. "In conclusion, detection of cocaine use from physiological measurements collected in the field setting is challenging," Kumar and Ertin, together with the student (lead) authors, wrote in the paper describing the work.

Sifting Through the Data Dump

For now, most mobile medical devices are designed to measure a piece of data that doctors already know they want—and has already been correlated to clinical outcome. Doctors already knew that lung fluid helps predict congestive heart failure, for instance; now they just have a new way to track that. And psychologists already knew that less active patients are more likely to be depressed.

But one promise of mobile medicine—at least for basic researchers—is that the vast amounts of data being collected can help reveal connections that clinicians don't yet know about. It's mobile medicine as a discovery tool.

"There's definitely this thought that because we have all this data, eventually we'll learn something from it," says Sim.

It's what online giants like Facebook and Google have already been able to do in other realms by tracking the online activity of people—discover what makes one person more likely to click an ad for a restaurant and someone else more likely to click an ad for a movie. Doctors want to do the same thing for healthcare: Ask what it is about a person's data footprint that can predict whether they'll get sick or whether they'll respond to a treatment.

"We can take all of these data points and find trends in them," says Joe Kvedar of Harvard Medical School and the Center for Connected Health. "It's what other industries have already done with complex, ever-changing data feeds."

A Common Language

As the data pool grows, another challenge emerges: the need for data from different devices to be standardized and compared.

When Sim orders a potassium level on a patient, for example, she doesn't care whether GE or Siemens produced the machine that measures it. "The data has to be device-agnostic," she says. "I can't be dealing with what machine the number came from." Likewise, if she wants to get information on a patient's daily activity for the past month, Sim doesn't want the information from a Fitbit to be presented to her differently than the information from an iPhone.

"We can take all of these data points and find trends in them," says Joe Kvedar of Harvard Medical School and the Center for Connected Health. "It's what other industries have already done with complex, ever-changing data feeds."

For now, though, that data pool just needs to grow. Is there such a thing as "too much data?" Most of the researchers interviewed for this story answered that question with a resounding no.

"There's certainly no such thing as too much data right now," says Markowetz. "That's the paradigm shift that's happening. Right now, you collect data because you can and because data storage is cheap. And then at some point later, you can say 'I have all this data, I wonder if it can help me answer a new question?'"

As long as researchers can keep up with managing the data, Topol says, there will always be a place for it. "There is this concern about TMI, too much information," says Topol. "And that is something that can be preempted by really great algorithms and analytics that filter out the signal from the noise and get the critical elements out of the data."

But sufficient data for researchers, Sim points out, might be too much for consumers or their doctors. "For individual patients and clinicians, there most definitely is a problem of too much data," she says. "We're almost there already." One key to moving forward, she suggests, is targeting the right data pools to the right people—clinicians might not need to see everything that researchers see, for instance.

The priority for tech companies, Sim points out, is to find their niche, market their product, and keep their data proprietary. But for researchers and clinicians, data that's in different forms depending on where it's from isn't useful.

"All these people who are trying to get data from multiple sources, and across heterogeneous platforms are starting to see the value of standardization," says Sim, who is a co-founder and a principal scientist of the non-profit Open mHealth, which aims to design an open, common language for health data.

Open mHealth has already outlined an

initial set of suggestions for standardizing such mobile sensor data, and Sim says they've garnered some interest from tech companies who want to learn more.

"People see the need for an open standard," she says. "We're only three years in and we're already getting traction. We'll be rolling out more this year."

The Tech World Meets Healthcare

Many of the mobile medical devices discussed here promise to give individuals more power over their own health. But mobile medicine also promises to ease the burden on our healthcare system by giving doctors tools to be more efficient, helping clinicians make the move toward truly personalized medicine, and providing platforms for clinical trials. But for this to happen, doctors and nurses will have to buy in to the power of wearable sensors.

"Doctors are not particularly enamored by many of these tools because it's a challenge to their control which has been in existence since the beginning of the profession," says Topol.

Changing their minds will require technology that's easy to use. "If you go into a doctor's office with a huge stream of EKG data that's been collected 24 hours a day, there's not much they can do with it," says Ertin. Software is needed to summarize these huge streams of data into a form that clinicians will find useful.

Yeung, as he moves forward with MyHeart Counts, imagines an algorithm that summarizes how each patient's heart health is doing—a green light means they're doing well, a yellow light might warrant a call from a nurse, and a red light would indicate the need for a virtual doctor's visit to discuss changing medications or approaches.

"This kind of system needs to be created to take the workload away from the physician," Yeung explains. "I could easily be following ten thousand or more patients anywhere in the world and only seeing the ones that need to be seen, rather than following fewer patients and scheduling frequent visits with all of them."

For the Menthal apps, Markowetz always aims to get a single number that captures a patient's status. "What if I had a single number that's your Parkinson's severity number per day?" he asks. "Now, I can chart this over the past six months, and that's something that really tells a doctor how you've been."

One day, healthcare may move this way. Rather than wait for your annual exam to have a doctor test the status of your health, you get daily updates, reminders, and notifications about your own body and behavior.

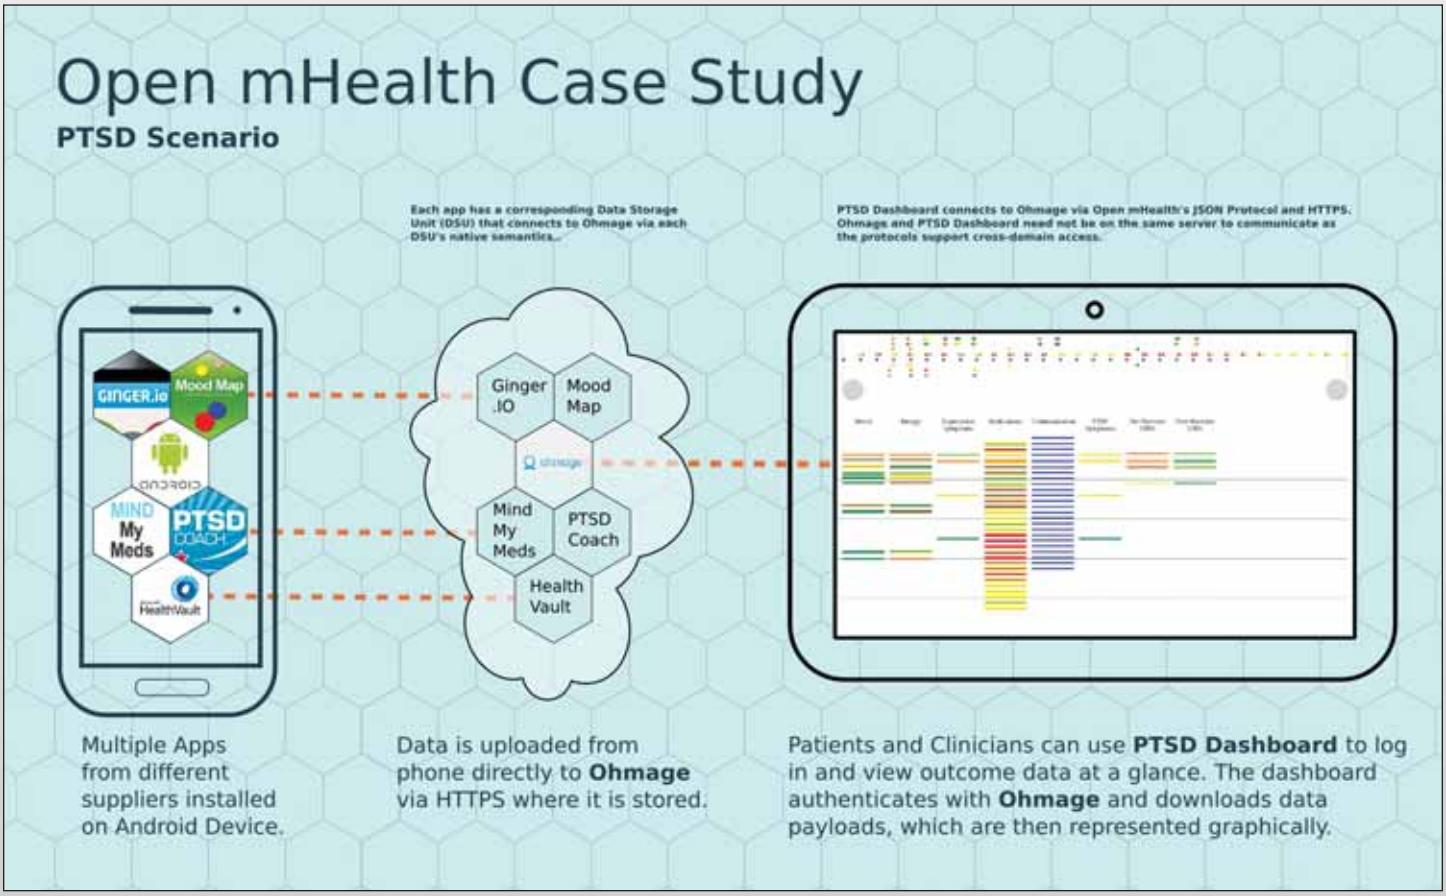

Open mHealth, in partnership with a physician, used a set of mobile apps on an Android phone to help a patient with post-traumatic stress disorder. The project helped the physician track symptoms, better understand the patient's condition, and intervene appropriately. Reprinted from http://www.openmhealth.org/open-mhealth-case-study-ptsd/

"This is going to be a whole rebooting of how medicine will be practiced to benefit the consumer," Topol says.

When it's needed, you'll be alerted to contact your doctor who will have a plethora of information at his or her fingertips to diagnose and treat you.

"This is going to be a whole rebooting of how medicine will be practiced to benefit the consumer," Topol says. □

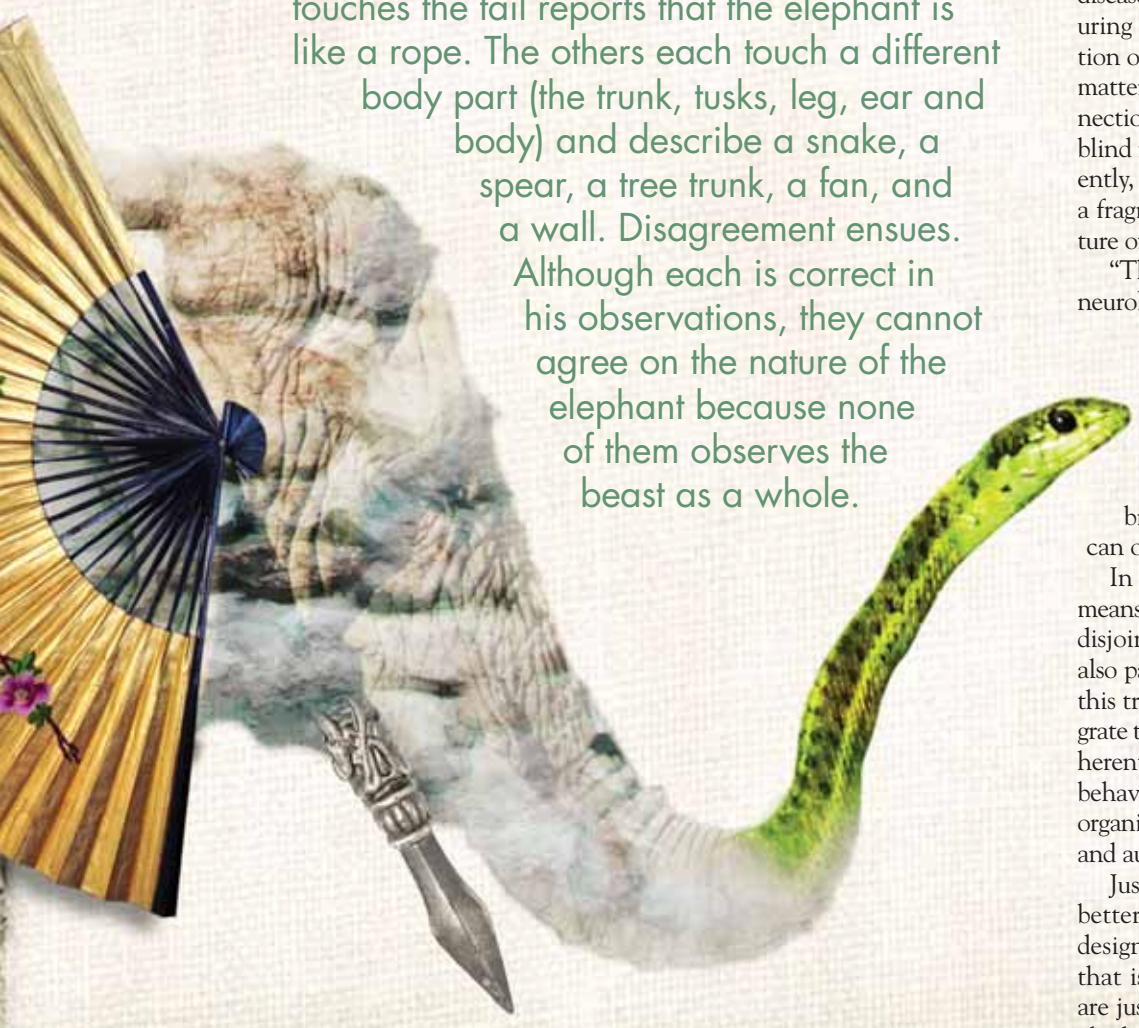

In an oft-cited story, six blind men each touch an elephant to describe its essential nature. The one who touches the tail reports that the elephant is like a rope. The others each touch a different body part (the trunk, tusks, leg, ear and body) and describe a snake, a spear, a tree trunk, a fan, and a wall. Disagreement ensues. Although each is correct in his observations, they cannot agree on the nature of the elephant because none of them observes the beast as a whole.

Integrating the FRAGMENTED MIND

Bringing the Whole Elephant into View

By Katharine Miller

The diverse researchers who study neurological diseases or psychiatric disorders such as schizophrenia are faced with a similar dilemma. They use multifarious approaches to understand the causes and effects of these diseases—sorting through the genome, measuring changes in the volume or concentration of gray matter, tracing the brain's white matter wiring, and spotting functional connections across brain regions. And just as the blind men each describe the elephant differently, so too do these various scientists report a fragmented and somewhat confusing picture of how mental illness affects the brain.

"The many approaches to understanding neurological disease and psychiatric disorders each offer a particular window into something that's gone awry," says **Arthur Toga, PhD**, professor at the Keck School of Medicine at the University of Southern California. But because all aspects of the brain work in concert, no single window can offer an integrated understanding.

In the case of schizophrenia, which means, quite literally, "fragmented mind," the disjointed nature of the research enterprise also parallels the disorder itself. People with this tragic mental illness don't seem to integrate their experiences of the world into a coherent thought process. As a result, they may behave in socially abnormal ways, have disorganized thoughts, and experience delusions and auditory hallucinations.

Just as their fragmented brains need to be better integrated, so too does the research designed to understand those brains. And that is now starting to happen. "Scientists are just beginning to join hands around the elephant," says **Olaf Sporns, PhD**, distin-

guished professor in the department of psychological and brain sciences at Indiana University, Bloomington. "They are collaborating more and also looking at the problem in all its complexity."

Some researchers are integrating structural information about the brain with genetic or functional data. Others tie genetics to phenotype or function. Still others are reaching for the whole enchilada using integrative systems analysis. While some pieces of the picture—such as environmental influences on the genome—remain out of focus, as the National Institutes of Health (NIH) and others are getting more interested in data mashing, progress is being made. "It's definitely the way to go," Toga says. "I think we'll see an accelerated pace of discovery because of it."

Genes and Schizophrenia

There's plenty of evidence that schizophrenia is highly heritable, yet no single genetic variant is the cause. And sample sizes have limited the productivity of genome wide association studies (GWAS). To address that problem, the Schizophrenia Working Group of the Psychiatric Genomics Consortium pulled together GWAS data from multiple institutions—amassing data for more than 36,000 cases and 113,000 controls. The study, published in *Nature* in July 2014, identified at least 108 genomic loci of significance. Many variants are located next to genes that operate in the brain or immune system—suggesting a possible link between the immune system and schizophrenia.

A separate study, also published in *Nature* in 2014, focused on identifying rare variants associated with schizophrenia by sequencing the exomes of 2,536 patients with schizophrenia and 2,543 unrelated controls. Individuals with schizophrenia had a significantly higher rate of rare disruptive mutations in protein-coding genes that were loosely suspected to play a role in schizophrenia. Moreover, disruptive mutations in 28 genes related to synaptic activity appeared in 9 cases versus none in controls; and disruptive mutations in 26 genes involved in calcium ion channels were found in 12 cases versus only one in controls.

Genes in these two gene sets appear to explain about one percent of schizophrenia cases. "It's consistent with the idea that there are many rare variants scattered throughout the genome, some of which probably confer risk for schizophrenia," says **Benjamin Neale, PhD**, assistant professor in the Analytic and Translational Genetics Unit at Massachusetts General Hospital, and an associated researcher at the Broad Institute.

Schizophrenia and the Thinking Brain:

A Matter of Gray Matter

Schizophrenia patients often experience disorganized thinking, hallucinations, and problems with attention, memory and language. And that in turn suggests a problem with the brain's gray matter. "The gray matter is where the good stuff happens," says **Jessica Turner, PhD**, associate professor of psychology and neuroscience at Georgia State University. "It's where the synapses are—where the cells fire. Without gray matter you can't think."

In imaging studies, researchers have long observed gray matter abnormalities in the brain structures of schizophrenic patients. Some such studies look at gray matter volume—the amount of gray matter inside the borders of particular brain structures—and some look at gray matter concentration—the density of gray matter judged by characteristics of the voxels. Most look for clusters of voxels that differ from healthy controls (univariate approaches) while a few have begun looking for patterns of variation

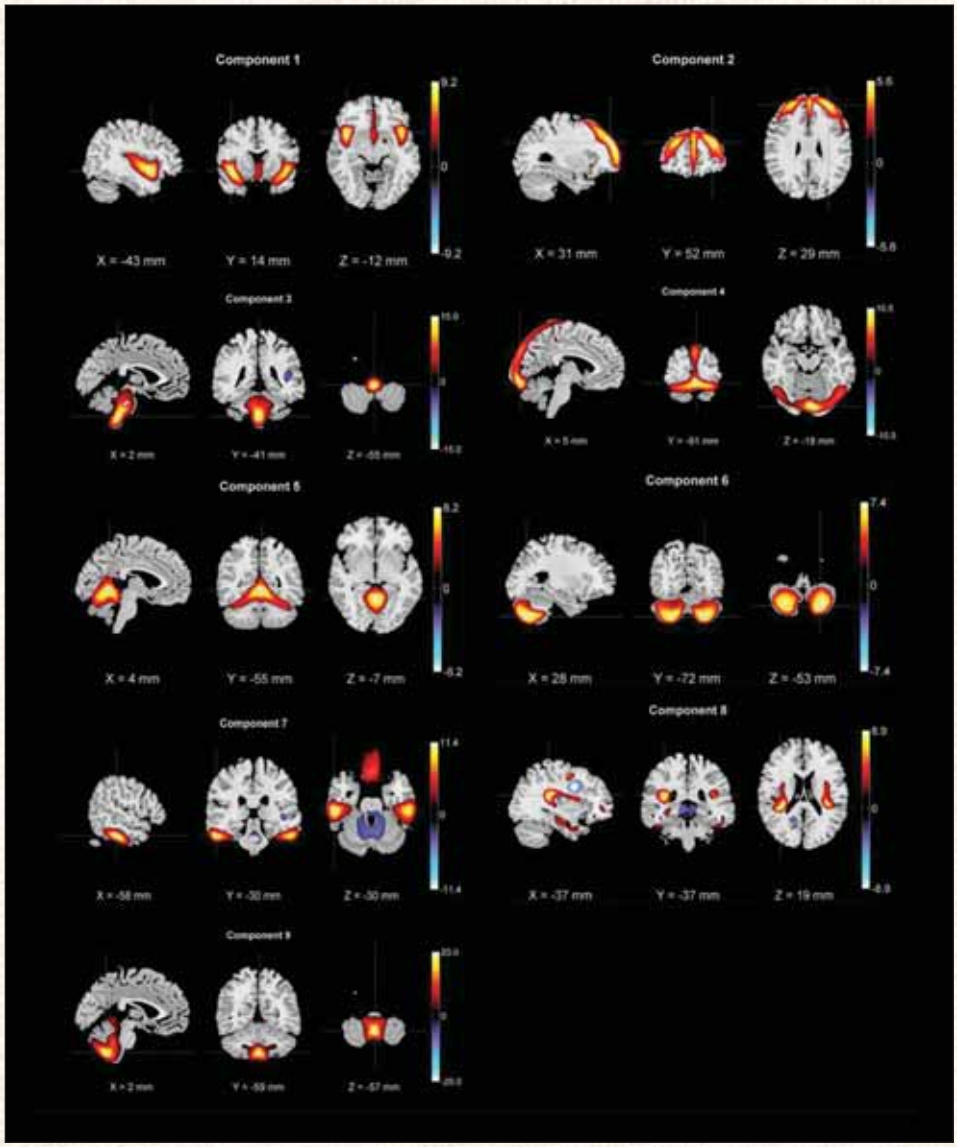

Turner's mega-analysis identified nine different spatial patterns (components) where gray matter concentrations in schizophrenia patients differed significantly from controls. Here, the nine patterns are shown in order from most significant to least. For the first seven components, patients had less gray matter than controls. Note that the spatial patterns were not defined by a brain atlas but rather revealed by the analysis of voxels. Red areas represent voxels where the differences were more highly statistically significant (z score > 2.5). Reprinted with permission from CN Gupta, VD Calhoun et al., Patterns of Gray Matter Abnormalities in Schizophrenia Based on an International Mega-analysis,* Schizophrenia Bulletin *(2014) doi: 10.1093/schbul/sbu177.

among those clusters (using multivariate approaches). Schizophrenia researchers also struggle with sample sizes that may be insufficient to reach statistical significance.

Seeking to generate reliable, replicable results, Turner and her colleagues set out to conduct an international mega-analysis of gray matter concentration using a multivariate approach and a large dataset. They gathered together MRI images from eight prior studies, including scans of 936 healthy controls and 784 people diagnosed with schizophrenia. The method they used—called parallel independent components analysis (ICA)—lets the data speak for itself, revealing spatial patterns that co-occur in patients compared to healthy controls. And the data did speak, revealing nine "components" or spatial patterns of interest. One pattern in particular caught the researchers' attention: reduced gray matter concentration in three areas (superior temporal gyrus, inferior frontal gyrus and insula) of the brains of schizophrenic patients. And the pattern was highly replicable. "You can find this over and over again in chronic schizophrenia," Turner says. "You can put it in the bank."

Is there potential for similar mega-analyses to reveal patterns in other mental illnesses? "Oh my goodness, yes!" says Turner. But she also points to the ENIGMA consortium as a model for future work. Rather than a mega-analysis, which brings all the data to one lab, ENIGMA leaves the data where it is and sends software scripts to participating investigators whose results are then combined. "ENIGMA gets more power out of collaboration and cooperation than we could out of doing our own little studies," Turner says.

For example, in 2008, the dynamic wave of gray matter loss that occurs as schizophrenia develops was revealed in a collaborative effort by 40 labs around the world led by **Paul Thompson, PhD**, professor in the Keck School of Medicine at the University of Southern California and director of the ENIGMA Consortium. In work published in 2008, Thompson and his colleagues also used time-lapse imaging to study the effects of various schizophrenia medications on the brain over the course of a year. Remarkably, they found one medication, olanzapine, that seemed to reduce gray matter loss compared with others.

Gray Matter and Genetics

Turner and her colleagues are now combining imaging and genetics approaches to schizophrenia to see if there's a relationship between the pattern of gray matter loss they observed in their mega-analysis and genetics. "Let's see if there's a relationship between this imaging pattern and cases/controls in GWAS," she says. An earlier project with a smaller number of subjects suggested the pattern of loss in schizophrenic patients is heritable. The results of Turner's team's genetics work are due out soon.

ENIGMA, which stands for Enhancing NeuroImaging Genetics through Meta-Analysis, is also driving forward with the emerging field of neuroimaging genomics. The idea is to use brain images to screen the genome for common variants that might affect the brain. To date, such approaches have identified genes linked to such things as brain or hippocampal size as well as a few genes linked to Alzheimer's Disease.

These early studies suggest, however, that neuroimaging sample sizes will need to be quite large if they are to avoid the problem of false positives—genes that seem linked to imaging features in a particular sample but cannot be replicated in other samples. This concern has led to the ENIGMA Consortium's efforts to combine images from many labs.

To date, the ENIGMA Consortium's Schizophrenia and Bipolar Working Groups have been focused on extracting meaningful information from neuroimages. But with support from the NIH's Big Data to Knowledge (BD2K) program, they hope to soon publish work that ties these results to genetics.

In a *Nature* paper published in January 2015, the ENIGMA network showed the potential promise of such work. That study found eight common variants (SNPs) that consistently predict the size of structures on brain MRI scans from over 30,000 people from 33 countries worldwide. Although non-genetic factors are clearly important, a tantalizing question is whether these genetic variants that correlate with brain structure also correlate with risk for brain diseases. For example, those eight genetic hotspots seem to affect the size of several brain regions implicated in schizophrenia, and some of them appear to affect risk for Alzheimer's disease, Parkinson's disease, and obsessive-compulsive disorder (OCD). An ongoing partnership between the Psychiatric Genomics Consortium and ENIGMA is comparing the two groups' find-

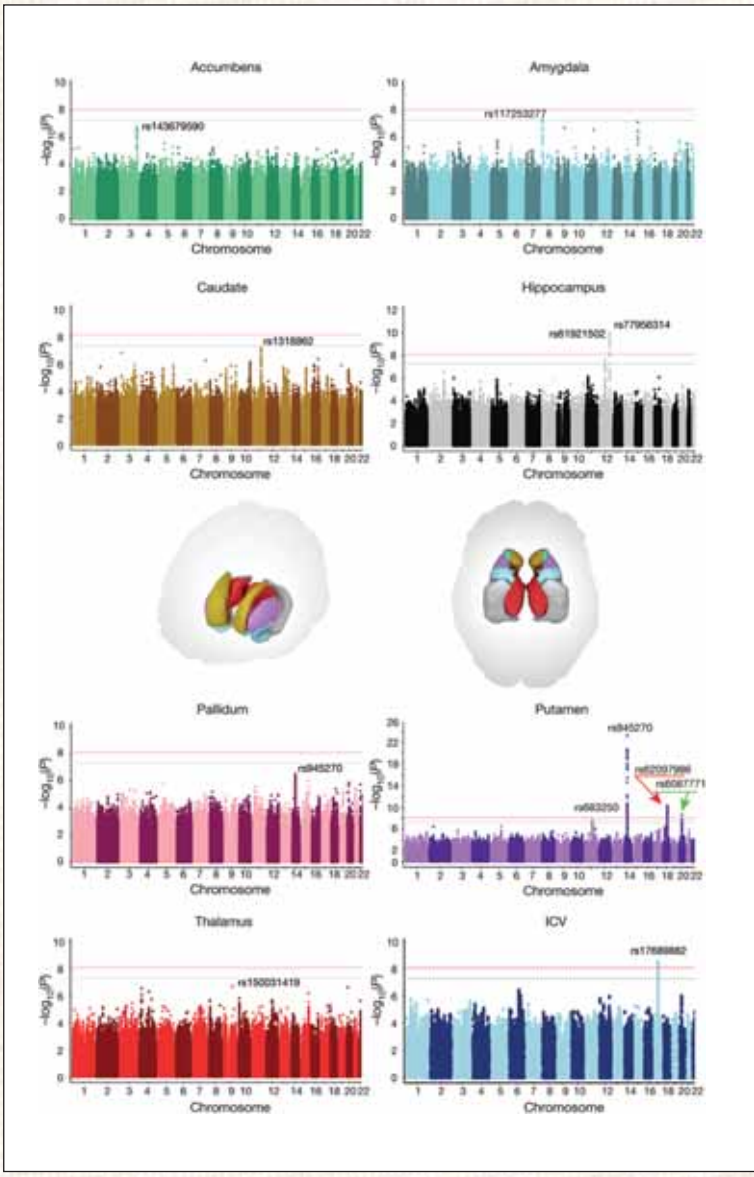

Recent neuroimaging genetics research by ENIGMA identified genetic variants (SNPs) associated with volume differences in various parts of the brain. These Manhattan plots are colored with a scheme to match the corresponding structure in the central diagram. Two different measures of genome-wide significance are shown with a gray dotted line ($P = 5\times10^{-8}$) and a red dotted line ($P = 7.1\times10^{-9}$). The most significant SNP within an associated locus is labeled. Reprinted with permission from Macmillan Publishers, Ltd., DP Hibar et al., Common genetic variants influence human subcortical brain structures,* Nature *(2015). doi:10.1038/nature14101.

ings to see if risk genes for schizophrenia might exert their affect by influencing the composition and integrity of the brain.

Schizophrenia and the Networked Brain:

A Matter of White Matter

According to a separate theory, differences in the brain's wiring could increase vulnerability to schizophrenia. In the brain, wiring means white matter—the bundles of axons that connect distant regions of the brain to one another.

Researchers can create wiring diagrams of the brain by tracing the diffusion of water along neuronal bundles, a method known as diffusion tensor imaging (DTI). And researchers like Sporns can then analyze these static images as networks. This approach has revealed some interesting things. For example, Sporns and his colleagues have found that highly connected parts of the human brain are also highly connected to each other, a characteristic called a "rich club." And, intriguingly, in schizophrenic patients, the connections between the members of the rich club are somewhat impaired while connections among less highly connected nodes are not.

Sporns thinks the rich club nature of the structural connectome is key to the brain's ability to function coherently. Our brains are constantly interacting with our environment and integrating information from many sources—our senses, memories, muscles, skills, internal physical states—to make sense of the world and guide our behavior in an integrated fashion. "Rich club, with its distributed pattern of highly connected hub nodes is analogous to a highway system for accomplishing this integrative task," Sporns says. But if a pathological mechanism weakens or disturbs that rich club connectivity, there's a penalty that is expressed in brain disorders such as schizophrenia, he proposes.

Bringing Functional Networks to Structural Networks

Having uncovered the brain's rich club structural network, Sporns decided to explore the relationship between static anatomical networks and functional networks that are much more dynamic, with changes on the scale of seconds or faster.

When brain researchers talk about function, they usually mean either how electrical activity changes among electrodes placed in the brain during an electro-encephalogram (EEG); or how blood flow in the brain changes over time (while at rest or doing a specific activity) as measured using functional magnetic resonance imaging (fMRI). "These methods produce a time series of neural activity using electrodes or voxels," Sporns says. "There's no movie that directly shows how neurons send messages to each other." So when researchers talk about functional connectivity in the brain, they are referring to activation patterns cross-correlated among different brain regions.

Just as network approaches can help researchers understand the structural connectome, so too can they reveal interesting features of the functional network. Previous work had shown that the functional networks of schizophrenia patients had reduced global communication capacity. Could that be due to reductions in rich club density? To find out, Sporns and his colleagues looked at structural and functional connectivity in the same patient population. And they found an increased coupling between the two types of connectivity in schizophrenia patients. That is, functional interactions were more directly related to the brain's anatomical connectivity in patients than in controls—possibly indicating less flexible or dynamic brain function in patients.

"Network approaches have given us a way of looking at schizophrenia from a different vantage point than we're used to," Sporns says. "And they've given us some hypotheses that we can now go out and test."

The Chronectome: Dynamic Connectivity

In a functional MRI, various parts of the brain light up together or separately in patterns that change through time. In the past, most fMRI studies have evaluated differences between patients and controls by essentially averaging these patterns during a particular activity and time period. These averages do reveal differences between people with and without schizophrenia. "But two brain regions might be highly correlated

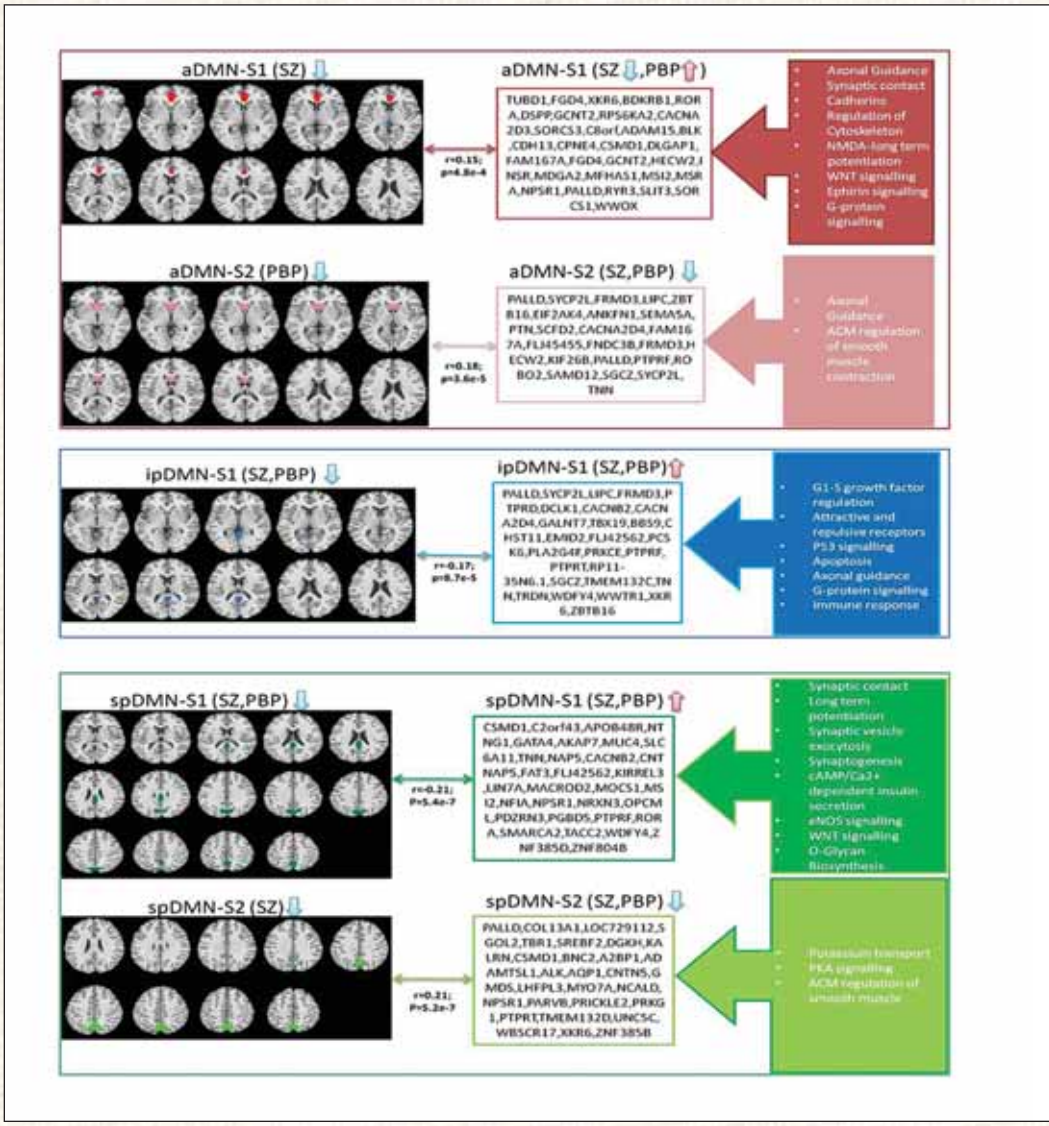

in the first few seconds and decline in correlation immediately," says **Vince Calhoun, PhD**, executive science officer at the Mind Research Network and distinguished professor of electrical and computer engineering at the University of New Mexico. "If you just look at average connectivity, you'll miss the change."

So Calhoun and his colleagues decided to scan people in a resting state to specifically look at whether connectivity dynamics themselves might reveal patterns of brain function that differ across people with schizophrenia or bipolar disorder and healthy controls. The work, discussed in *Neuron* in November 2014, found five states (correlations among specific regions) that exist routinely in both cases and controls, but when they looked at the "dwell time"—the percent of time spent in each state—the schizophrenia patients occupied two particular states much longer than the controls. "You do learn some interesting things from the averages," Calhoun says. "But when you unpack it and look at what goes into that average, you actually learn a lot more."

Since Calhoun's first dynamic connectivity paper was published in 2009, the approach has "kind of exploded," he says. "It reinforces the value in looking at the data that way."

Integrating Genetics and Phenotypes

Standard GWAS can only point to individual variants that are associated with a mental disorder, says **Igor Zwir, PhD**, postdoctoral scholar working with **Robert Cloninger, PhD**, at Washington University in St. Louis. And many of the findings are weak and inconsistent. In any event, Zwir says, "Individual genes basically do not cause mental disorder. Genes act in concert as well as with the environment."

Moreover, because schizophrenia is actually a spectrum of disorders that varies widely in severity and covers a whole range of symptoms from positive (such as delusions, disordered thoughts) to negative (such as lack of interest in others, inability to feel pleasure or act spontaneously), Zwir notes, the gene clusters that interact to produce different sets of symptoms may be different as well.

So Zwir and his colleagues decided to take a self-organizing approach to a large set of GWAS data for about 4,000 schizophrenia patients whose symptoms and their severity were also well documented and a similar number of healthy controls. Without any presumption as to which gene mutations (SNPs) would co-occur, they partitioned the patients based on shared sets of SNPs. They did the same thing for phenotypes—allowing the data to cluster people without any presumption as to which traits go together.

Next, they optimized the relationships between the clusters of SNPs and the clusters of traits. And for each association they calculated the risk. "If there's a 90 percent risk, then that association includes 90 percent cases and 10 percent controls," Zwir says. Ultimately, the clustering honed in on eight sets of SNPs with associated phenotypes.

When Zwir and his team reported these results in September 2014 in the *American Journal of Psychiatry*, the press coverage was intense, with many publications declaring the existence of eight subtypes of schizophrenia. The researchers received calls from *Time*, *Newsweek* and even Anderson Cooper. "Why did we have this press?" Zwir asks. "It's because people need to relate genetics with disease and it isn't often done."

At the same time, some members of the research community criticized the group's methodology. Without weighing in on the details of that debate, Turner commented that no matter how one feels about the statistical details, "This is a very rich, very reasonable approach, and the findings made sense."

Moreover, because of the Zwir paper, Turner is now trying to apply similar methods to imaging data. But it's difficult to get phenotypic data that's properly standardized, she says. "When you try to break it down and look at what exactly the symptoms are and how bad the hallucinations are or how bad the reality disorganization or cognitive deficit, they are not well quantified. Different people use different scales."

Undeterred by the critics, Zwir and his colleagues have recently applied the same clustering algorithm to see whether schizophrenia phenotypes cluster with different patterns of white matter loss (using DTI) in schizophrenia. Though their sample size was relatively small (47 patients and 36 healthy controls) they found at least three distinct clusters of symptoms and white matter patterns—one pattern associated with bizarre behavior; another with prominent delusions; and a third with negative symptoms, including disorganized speech. The work was recently submitted for publication.

Integrating Genetics and Function

In another effort to integrate different perspectives, Calhoun worked on a project to find genetic patterns that coincide with brain functional network patterns in a group of subjects that included patients with schizophrenia and bipolar disorder as well as their healthy family members and healthy but unrelated controls.

The work, under the leadership of **Godfrey Pearlson, MD**, at Yale, used a large array of genetic information (single nucleotide polymorphisms) and focused on the default-mode network (DMN) of the brain. "It tends to be more active when you're not focused externally," Calhoun says. This network typically shows reduced functional connectivity in people with schizophrenia and bipolar disorder. Using parallel independent component analysis, an approach developed by Calhoun's group, the team was able to find genetic patterns and DMN subnetwork patterns that co-occur in a group of subjects. "This gives us a richer set of features that we can pull out of the data without starting from a region of interest," Calhoun says. The group then went further, and sought to understand the possible molecular underpinnings of the genes identified in the study. The results, published in

(Opposite page) Calhoun and his colleagues parsed a dataset of functional scans and genetic information for schizophrenia and bipolar patients using a method Calhoun's lab developed called parallel independent component analysis. The technique revealed five subnetworks of the brain's default mode network (the part that's focused internally) and their associated genetic components, represented here by the top 10 genes as well as genes that feature more than three SNPs within each network. The blocks at right list significantly enriched ontology terms within each genetic cluster. Arrows pointing up and down indicate whether the loading coefficient for that particular feature (fMRI or gene) was significantly higher or lower for patients or controls. Reprinted with permission from SA Meda et al., Multivariate analysis reveals genetic associations of the resting default mode network in psychotic bipolar disorder and schizophrenia,* PNAS *vol. 111:19 (2014).

PNAS in April of 2014, pointed to mechanisms that had been previously implicated in schizophrenia and bipolar disorder as well as several novel mechanisms.

Integrative Systems Analysis

Michael Snyder, PhD, professor of genetics at Stanford, has long been in the business of data integration. Since the 1980s he has been working with combinations of data including genetics, genomics, transcriptomics and metabolomics. "We're pretty comfortable working across these areas and integrating lots of different information," he says.

Snyder and his colleagues took the entire human protein interaction network (the "interactome") and mapped its organization at an intermediate scale. "If you imagine all the proteins are the world, the map we set up is kind of like at the state level, where groups of proteins are working together," he says. They then took known genes for autism spectrum disorder (ASD) and mapped them onto these clusters. "Two modules screamed out at us," Snyder says. "But especially one. Autism kids have a high chance of mutations in our module."

Snyder's team didn't stop there. They used whole genome sequencing to look at 25 kids with autism and found they were enriched for mutations of genes in the module. They then looked at what the module does using the Allen Brain Atlas, and found that half are expressed in most neurons but half of them are primarily expressed in the corpus callosum. The importance of the corpus callosum in ASD was also confirmed in mouse models.

"I think this kind of analysis will be very fruitful when applied to other areas," Snyder says.

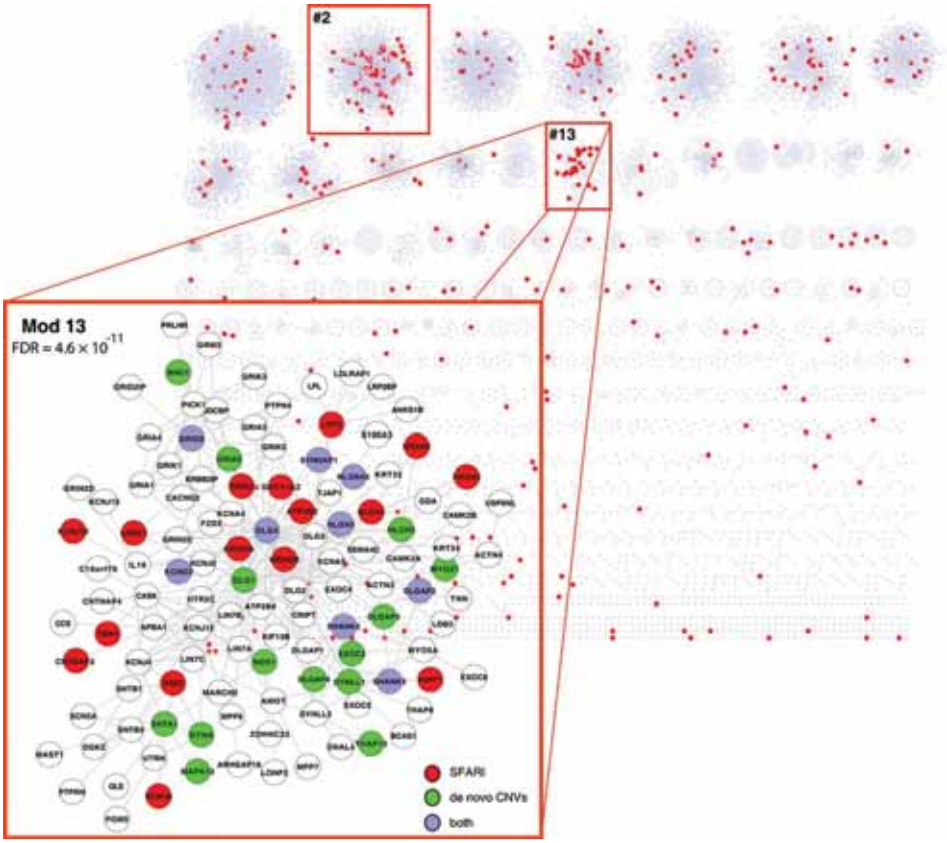

Snyder and his colleagues identified genetic modules (#2 and #13 above) in the human protein interactome that are enriched for autism–associated genes (in red). The topological modules are physical clusters on the protein interaction network where member genes intensively interact with each other but sparsely interact with non–member genes on the network. The zoom–in view of module #13 is colorized to show known autism genes (red) and genes affected by autism spectrum disorder–associated* de novo *copy-number variations (green). Genes annotated by both were in blue. Reprinted from J Li, M Shi, Z Ma, S Zhao, et al., Integrated systems analysis reveals a molecular network underlying autism spectrum disorders,* Molecular Systems Biology, *10 (12) 2014.

For diseases like schizophrenia and other psychoses, such a combination could be quite powerful, Snyder says. He bases that assessment on recent work his team did in the area of autism.

Data Fusion: Bringing Multiple Imaging Approaches Together

Calhoun is very interested in pulling together multiple types of imaging to see what they can show us about the brain. Often, researchers integrate imaging data by overlaying one image on another. However, such approaches will not necessarily recognize if a change in one area of the brain correlates with a change in another part of the brain.

Calhoun favors a different approach he calls data fusion, in which each method informs the other without any assumptions about which information is more important. "We don't, at the beginning, make a critical assumption that might lead us down a wrong path," he says. "My approach is to move the simplification step to the end."

For example, at the IEEE Engineering in Medicine and Biology Society Annual International Conference in 2014, Calhoun and his colleagues presented a data fusion approach to combine three types of imaging data—functional MRI, EEG and structural MRI (white and gray matter volumes)—in a study comparing schizophrenia patients and controls. They found that the combined data was more predictive of schizophrenia status than any single imaging modality alone.

Treating the Elephant

As researchers begin to grasp hands around the elephant, they may start to find connections that explain how genetics and the protein interactome play into gray matter loss and reduced rich club network structure across the full range of psychotic phenotypes.

Ultimately, such an integrative approach could lead to better diagnosis and treatment options for the millions who suffer from schizophrenia or other psychoses.

Turner looks forward to the day when a brain scan and genetic test can help physicians steer people toward appropriate treatments based on a molecular understanding of what's going on; or allow early interventions in young people, to prevent gray matter loss before it can get started. To get there, many perspectives on the psychotic brain will have to collectively tell one story. □

Under TheHood

BY REZA BOSAGH ZADEH, PhD

Machine Learning using Big Data: How Apache Spark Can Help

Machine learning is the process of automatically building models from data. In the past two decades, researchers in many fields of study have been generating these models from progressively more data. Because this has led to higher quality learned models, researchers are using even greater quantities of data that require more and more complex distributed computing systems. These systems consist of many hard-drives connected to many machines (CPUs)—often commodity computers to keep costs down. But with many commodity machines come many failures: Hard-drives die; operating systems fail; and someone might trip over a power cord in the data center. The need to problem-solve such single points of failure renders distributed computing quite cumbersome. One solution: Use cleverly designed software to make applications running in clusters more fault-tolerant. Specifically, researchers turn to software known as cluster programming frameworks.

The most successful of these is Apache Spark. Built by the AMPLab at the University of California, Berkeley, and now controlled by Databricks, Spark provides users with a distributed array that is fault-tolerant. Many researchers are already accustomed to programming with arrays in their favorite programming language. Spark provides much of the same functionality that arrays provide, with the convenience of the array being seamlessly distributed across a cluster. These arrays are called Resilient Distributed Datasets (RDDs). They can be large and stored on disk, with the portions that are in use swapped in and out of RAM for faster access. Because the generic idea of distributed arrays has nothing to do with any particular programming language, Spark is able to provide clean APIs in Python, Java, Scala, and R.

There are many ways to create RDDs, but the world only lets you create RDDs in ways that can be automatically tracked. The recipe for an RDD is saved along with the RDD, so that in the event of a machine failure, the part for which the machine was responsible can be rebuilt. Called "lineage," this recipe is the primary fault-tolerance mechanism in Spark.

Many researchers are already accustomed to programming with arrays in their favorite programming language. Spark provides much of the same functionality that arrays provide, with the convenience of the array being seamlessly distributed across a cluster.

Given that programming with arrays has been historically successful, it is no surprise that RDDs have also enjoyed fast adoption. Spark provides four libraries out of the box that take advantage of the power of RDDs:

- **ML**: Machine learning algorithms and matrix computations
- **GraphX**: Graph processing library for handling large graphs
- **Streaming**: Handling streams of data (e.g., web logs or stock tickers)
- **Dataframes**: Easy access to tables of heterogeneous data, similar to those found in R and Python

These open-source libraries are developed in a concerted effort across many universities and companies. For example, several Stanford students have worked with me to create the basic building blocks for linear algebra in Spark, such as the singular value decomposition. Only the most widely used and tested algorithms are added to the above libraries. However, there is a vibrant community of people developing Spark packages that can be installed with a single command line. Databricks maintains this package listing at http://spark-packages.org. Together, the Spark ecosystem and its community make big data easier to handle. □

DETAILS

Reza Bosagh Zadeh is a Consulting Professor at Stanford University and a Technical Advisor to Databricks. Zadeh received his PhD in Computational Mathematics from Stanford University under the supervision of Gunnar Carlsson. For his PhD work in distributed machine learning, he received the Gene Golub Outstanding Thesis Award. During his PhD, Zadeh built the machine learning algorithms behind Twitter's who-to-follow system, the first product to use machine learning at Twitter. As a Technical Advisor at Databricks, Zadeh is the initial creator of the linear algebra package in Apache Spark.

DIVERSE DISCIPLINES, ONE COMMUNITY
BiomedicalComputation REVIEW
Published by the Mobilize Center, an NIH Big Data to Knowledge Center of Excellence

Nonprofit Org.
U.S. Postage Paid
Permit No. 28
Palo Alto, CA

Stanford University
318 Campus Drive
Clark Center Room S271
Stanford, CA 94305-5444

seeing science

SeeingScience

BY KATHARINE MILLER

Seeing Inside

The bright light needed to see molecular activity inside a living cell can quickly alter or even halt the very thing scientists want to observe. But a new technique developed by **Eric Betzig, PhD**, Group Leader at the Janelia Research Campus, offers fantastic 3-D resolution of living cells for longer time periods without phototoxicity. Called lattice light-sheet microscopy, the technique uses ultrathin light sheets derived from two-dimensional optical lattices. Rapidly scanned plane-by-plane through the specimen, these light sheets provide excellent illumination with minimal damage to the cell. Betzig, who won the 2014 Nobel Prize for Chemistry for other work, calls lattice light-sheet microscopy "the high-water mark" of his career. □

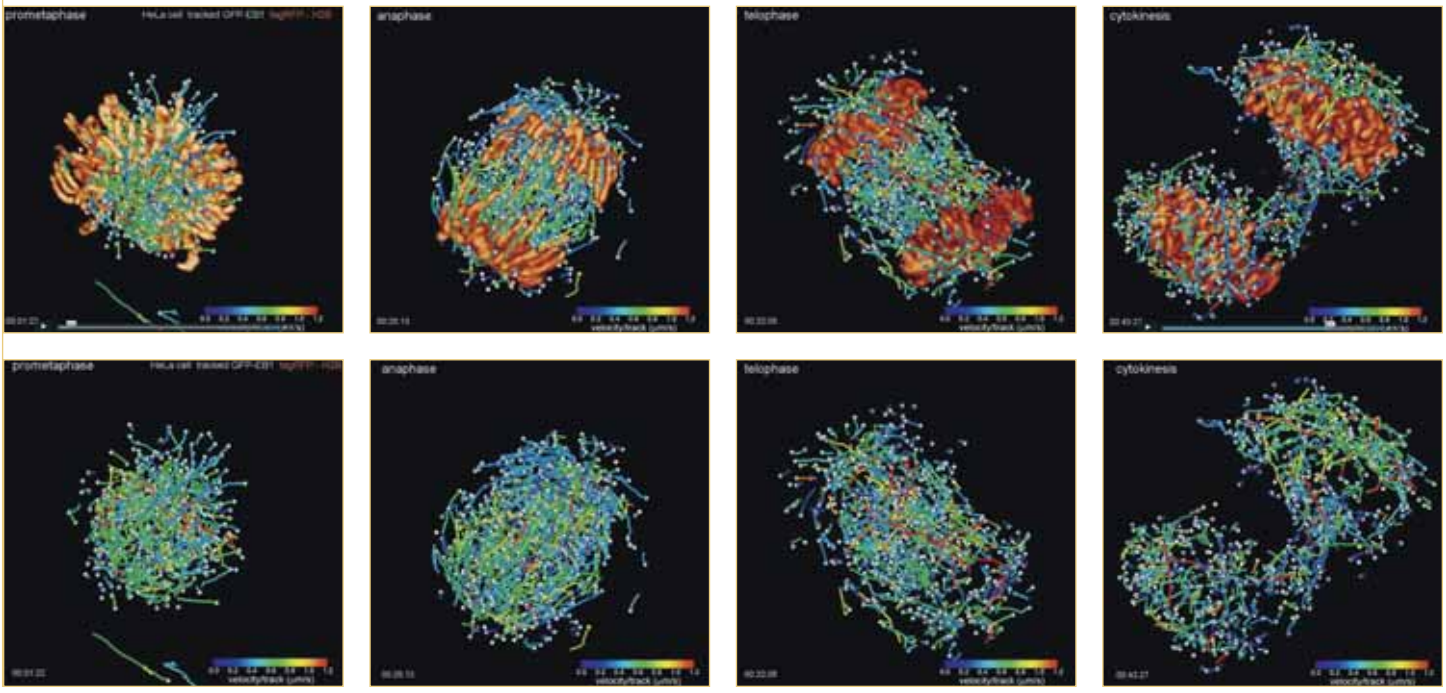

Lattice light-sheet microscopy allows the imaging of molecules inside living cells. Here, HeLa cells progress through mitosis with chromosomes (in orange) and the 3-D growth and retraction of microtubule components shown as points with lines colored according to their velocity. Credit: Betzig Lab, HHMI/Janelia Research Campus; Mimori-Kiyosue Lab, RIKEN Center for Developmental Biology. Reprinted with permission from B-C Chen et al., Lattice light-sheet microscopy: Imaging molecules to embryos at high spatiotemporal resolution,* Science *346:6208 (2014).